The Value of Independent Component Analysis in Identifying Climate Processes

Thesis presented for the degree of
Master of Science

Michael Kent
June 2011

Supervised by
Bruce Hewitson and Dáithí Stone

Climate Systems Analysis Group
Department of Environmental and Geographical Science
The University of Cape Town

Abstract

The variability of the atmosphere often displays distinct processes that allow us to better understand and predict it. As the processes are smaller and simpler than the atmosphere itself, these can be combined to better understand and predict the variability of the atmosphere as a whole. However theories about complex processes often lack supporting evidence. Therefore methods can be used to identify patterns in data which are characteristic of the processes. The successful identification of the patterns then serves as supporting evidence for the correct understanding of the processes. We use Independent Component Analysis as a method to identify patterns (signals) in data that are distinct from noise. We more specifically use the non-Gaussian signals that are statistically independent from each other to identify processes within data.

To determine how important a signal is to the identification of a process, the percentage of variance explained (PVE) by the signal can be used. The larger the PVE by the signal the more atmospheric variability it is able to predict, and thus the more important it may be. However, we demonstrate that the PVE by signals is sensitive to both the number of signals separated from the data and to the use of random values as the initial estimate of the signals. We find that only the first few signals that have large PVE values are insensitive to changes in their initial estimate. The remaining signals are all sensitive to the initial estimate change. Furthermore, there appears to be no relationship found to exist between the sensitivity of the remaining signals, and the PVE by them when changing their initial estimate.

Previous work in ICA has found that a single process can be represented by multiple signals. We introduce the use of spatial maps to aid in the linking of signals to processes. The spatial maps represent the spatial manifestations of a signal over time. They are used to strengthen the link between a signal and a process by seeking a spatial pattern within the spatial maps at the corresponding time of a known climate event (Eg: El Niño). If the spatial patterns are representative of the process at the corresponding time of the event, then the link between the signal and the process can be strengthened. We find that in our experiments, the land based seasonal cycle is represented by the signal with the largest PVE, while an ocean based seasonal cycle is generally identified using the signal with the second largest PVE value. We also demonstrate that the spatial maps were able to identify the North Atlantic Oscillation, but generally failed to detect the El Niño Southern Oscillation. An additional value of ICA is that it is able to detect data inhomogeneities, as we found that the change of NCEP to using full satellite data in 1979 was represented by two signals from different sets.

We find that ICA could make a useful contribution through the identification of the land based seasonal cycle and the ocean based seasonal cycle. These qualities mean that ICA may further prove a useful tool in the problem of identifying components of climate change signals from ensembles of multiple climate models.

Plagiarism Declaration

In compliance with the Postgraduate Handbook from the Faculty of Science:

"I know the meaning of plagiarism and declare that all of the work in the document, save for that which is properly acknowledged, is my own".

Acknowledgments

My supervisors, Bruce Hewitson and Dáithí Stone, for reviewing many drafts and for helping me identify results that I would have never seen on my own. Chandrika Kamath, for responding to a request for details on the implementation of the prototype. My neighbours, for finally feeding their cat. And lastly, my family and friends who planned around me and postponed many events in order to ensure that I could make it to them.

Contents

1 Understanding Atmospheric Variability

1.1 Introduction

The variability of the atmosphere often displays distinct features or processes that allow us to better understand and predict it. As each process represents a smaller simpler part of the atmospheric variability they give us an accurate description of its behavior when they are combined. Thus by understanding these simpler processes of the atmosphere we are able to better understand the atmosphere as a whole.

This thesis explores the question of identifying atmospheric processes by employing a new technique, Independent Component Analysis (ICA), to identify different signals of the climate system without prior knowledge of the signals themselves. In doing so the thesis assesses how valuable the method is in identifying patterns which are characteristic of processes, with a view to developing future work in evaluating climate model simulations.

The seasonal cycle of the atmosphere is a good example of a process. The seasonal changes of the atmosphere are clearly and easily identifiable. The North Atlantic Oscillation is another example of a process. It however is not as predictable as the seasonal cycle. The two pressure systems which define its behavior are dynamic and it is difficult to derive patterns from them. Although it is smaller than the seasonal cycle it is still important as it effects the regional climate. Because of the difficulties in trying to understand it, many ideas and concepts have been developed to tackle the problem from different perspectives.

But the many different perspectives are not all in agreement with each other as to how the North Atlantic Oscillation behaves. So in an attempt to reconcile some of the approaches and concepts, methods can be employed to identify the process within data. If the process is successfully identified, then the evidence to support the understanding of it grows. An example of this can be found in the work by Wallace and Gutzler (1981). They examine data using correlation values, and find patterns in them which are characteristic of the North Atlantic Oscillation. In doing so they identify the North Atlantic Oscillation as being represented within their data, which then further reinforces their understanding of it.

The example reveals the methodology of building support for the understanding of a process by the identification of patterns which are characteristic of it. In general this methodology is repeatedly used to both build support for the processes, and for the methods which identify them. As each new method employed to identify processes is different, they offer the potential to reveal new insights into the behavior of the processes which they identify. These insights can then further our understanding and predictability of the processes. It is in order to seek new insights that we propose the use of Independent Component Analysis as an additional method for identifying processes.

The difference with ICA is that it identifies patterns which are distinct from the pattern of noise. Noise normally represents a mixture of processes which by itself does not make much sense. However, with ICA the individual processes within the noise can be identified. This makes ICA able to uncover processes which previous techniques have been unable to.

1.2 Methods for Identifying Climate Processes

This section introduces some of the methods that are available to identify processes. We now replace the term process with the more specific term of climate process. We use the term c*limate process* to refer to a structurally cohesive atmospheric feature, that describes a part of the behavioral variations in the atmospheric dynamics. Examples of climate processes include the seasonal cycle and the North Atlantic Oscillation.

We review some of the existing methods, which focus on correlation, similarity measures, and decorrelation to identify climate processes. We then propose Independent Component Analysis as an additional method that separates variables which are statistically independent of each other and have non-Gaussian distributions. We show that the value of ICA lies in its ability to identify non-Gaussian variables in the data where other methods would fail to identify them. The different methods for identifying climate processes can be divided into the following.

1.2.1 Correlation Maps

The key work in this area is by Wallace and Gutzler (1981). They investigate correlations between grid points to identify spatial patterns which may be evidence of climate processes. To identify the spatial patterns they construct a correlation map for each of the grid points. A correlation map is created by correlating a grid point (the basis point) against all the other grid points in the data. The correlation values within a correlation map are geographically distributed over the region of the data, and are known as the *Teleconnections*. Therefore each correlation map is associated with a basis point and consists of Teleconnections.

To determine which basis points contain Teleconnections that are the most representative of climate processes, they examine the strongest negative Teleconnections in all the basis points. They reason that if strong climate processes are reflected in a correlation map, then the corresponding grid points effected by the processes will be dissimilar to the basis point and therefore be negatively correlated to it. They plot the maximum negative correlation values onto an image which represents the spatial coverage of the data. The location of each basis point in the image is set to the corresponding maximum negative correlation value found within its respective correlation map.

The regions in the image with high negative correlations (absolute values $\geqslant 0.75$) indicate basis points which contain strong negative Teleconnections. The basis points within the regions with the same highest

negative correlation value are known as the *centers* of the regions. As the centers have the strongest negative correlation values, they are deemed to contain the Teleconnection patterns which are most representative of climate processes.

In order to identify climate processes, they compare the Teleconnections of the centers to the spatial patterns of documented climate processes. They find that the Teleconnection patterns for both the geopotential height and SLP data are also present as spatial patterns in both the North Atlantic Oscillation and the Pacific/North American Pattern. In doing so they identify these climate processes within the same spatial region and time period as their data.

1.2.2 Communities in Graphs

As an alternate method to using correlations between grid points to identify climate processes, clustering methods employ similarity measures (eg: area-correlation) to find new groupings amongst the grid points of the data. These groupings are known as clusters or communities, and with geographically distributed data, the spatial extent of the communities can represent spatial patterns within the data. In clustering methods, a graph represents a series of nodes (Eg: grid points) which can be linked together into communities based on a measure of their similarity. The created communities can then be used to identify climate processes.

Steinbach et al. (2003) seek to uncover known climate indices (Eg: Niño 3.4) which represent the time series of climate processes. By finding the known indices they aim to establish a methodology with which to identify new indices. They apply the Shared Nearest Neighbour clustering algorithm to uncover climate indices within observational SLP and SST data. The similarity measure of the nodes is based on the area-weighted correlations of known climate indices to the data. The average of all the index correlations, creates a threshold which constrains the spatial extent of the communities. Additionally, the time series for a community is the average of all the nodes in the community, and is termed the *centroid* of the community. They find that the centroids have high correlations to the known climate indices, and that the communities can reveal new spatial patterns about the indices. As the formation of the communities are dependent upon the number of known climate indices used, a change in the number of indices used may lead to a change in the spatial patterns of the communities as well.

The WalkTrap algorithm is used by Steinhaeuser et al. (2009) to investigate the clustering of nodes when considering multiple variables[1] as opposed to a single variable. Using the absolute correlations between the variables associated with the nodes, the similarity between the nodes can be determined. The more similar the nodes, the smaller the distance between them. The distance between the nodes therefore serves as the similarity measure of the nodes. After creating a link between every node in the graph they retain only the top 1 percentile of the strongest links corresponding to the shortest distances between the nodes. The

[1] The variables: Air temperature, pressure, relative humidity, and precipitable water.

WalkTrap algorithm forms clusters from the retained nodes on the assumption that random walks can be used to calculate the distance between the nodes and clusters. A node can be merged into a cluster only if the distance between the node and cluster is sufficiently small. Their results from observational data show four global patterns representing climate processes.

As Christiansen (2007) points out, often clustering methods need the number of clusters to be known prior to application of the method. This can therefore make it difficult to apply them in any context in which the number of clusters are not known prior to the analysis. They also caution the use of clustering methods, due to their lack of robustness and because their statistical properties may not be well defined.

1.2.3 Principal Components

Principal Component Analysis (PCA) provides an alternative perspective as to how the climate processes are related to each other within data. Rather than focusing on correlation or a similarity measure, PCA produces decorrelated variables which are orthogonal to each other and can be used to identify climate processes. More specifically, PCA is a method that linearly rotates data into a set of decorrelated variables. The variables are known as Principal Components (PCs), and each PC is orthogonal to every other PC. The first PCs maximises the amount of variance that it explains, with each subsequent PC explaining less of the total variance of the data. (PCA is further discussed in section 2.3).

Compagnucci and Richman (2008) investigate what effects are seen in the results when using different modes of PCA. S-mode PCA assumes that the time series at each of the data grid points are the variables of interest. On the other hand, T-mode PCA assumes that the spatial field defined by the grid points at each point in time are the variables of interest. To investigate the effects of the different modes, they construct an artificial data set with known properties. Then by comparing the results obtained by using the different modes to the known properties, they are able to compare the effects of using the different modes. They find S-Mode PCA is best used for finding spatial clusters or teleconnections, while T-mode PCA is best used for finding spatial synoptic or flow patterns. Therefore the different modes constrain the conclusions that can be drawn from the results.

Barnston and Livezey (1987) state that one concern with PCA is that each subsequent PC is less robust and that this is particularly problematic when considering different periods of the data. They propose the Orthogonal Varimax Rotated PCA (RPCA) method to address this concern. Additional motivation behind RPCA is that simple structures can be found to represent climate processes by rotating principal components. The resulting simpler structures have only a few highly correlated values (positive or negative), with the rest being close to zero. These simpler structures are stated to be more representative of Teleconnections than the unrotated PCs. They apply RPCA to Northern Hemisphere 700mb height data, and find a series of spatial patterns within the Rotated PCs which represent spatial patterns of climate processes (eg: North Atlantic

Oscillation). To test how robust of the patterns are, they show the correlations between monthly results from the same pattern and the results from other patterns. The inter-pattern correlations are found to be weaker than the intra-pattern correlations, indicating that the pattens are robust and are not artifacts of the analysis procedure.

Dommenget and Latif (2002) outline the assumptions used in PCA and RPCA techniques. They argue that the assumptions about the climate processes and the techniques are different. They reason that the orthogonality constraint imposed by the techniques is not always consistent with climate processes, which can be non-orthogonal in space. Indeed, when using an artificial dataset, they find that the resulting patterns from each technique can contain a combination of climate processes, rather than an individual process. This makes the association of a pattern to a single climate process difficult. As a solution they promote the use of multiple techniques in order to make conclusions which are representative of the data, and are not artifacts of a single analysis technique.

The Local Model Analysis (LMA) method by Goulet and Duvel (2000), applies a series of Complex PCA steps to a region to examine a particular oscillation in terms of its composite structures. These structures can then be used to determine the most persistent underlying oscillation that is present within the data. Complex PCA augments PCA by additionally revealing spatial information on the different phases of the oscillation. In using the LMA method to apply Complex PCA to smaller time periods (windows) of the data, they remove the orthogonality constraint of PCA, which allows them to investigate any intermittent and non-orthogonal oscillations which may be present in the data.

A subset of the results from the LMA method are retained, and these represent the underlying structures within the oscillation. If the structures persist in multiple windows, then the structures are deemed to represent a persistent oscillation. They find that if the spatial domain is made larger than that of the investigated oscillation (Eg: Inter-seasonal Oscillation), then other oscillations are found to interfere with the analysis. In addition to this, they find that the length of the window may have to be made longer so as to identify structures which are more persistent, while avoiding the mixture of oscillations between windows.

Analysis methods such as PCA, show that climate processes can be mixed over space and time within data. Jolliffe (2003) reviews the work by Dommenget and Latif (2002) and proposes Independent Component Analysis (ICA) as a method for unmixing statistically independent variables from data. Using ICA, we intend to to find the unique behavior of process by imposing the constraint of independence between them. Although climate processes are not independent of each other in the truest sense, it may be assumed that they are independent in an attempt to extract variables which may better represent the theoretical understanding of climate processes, rather than mixtures of them. We assume this in order to find the unique behavior of each climate process.

Furthermore, we examine the climate processes as being distinct from Gaussian noise, in other words maximising their non-Gaussian distributions using ICA. Good motivation for the presence non-Gaussian variables in data is provided by Aires et al. (2000) where they apply ICA to SST data. They show that within their SST data, the El Niño Southern Oscillation can be identified using the non-Gaussian results of ICA. This suggests that climate data can contain non-Gaussian variables which can be found using ICA, while PCA and RPCA would fail to identify them. Therefore, an understanding of ICA will provide a complementary analysis method for identifying climate processes in data. In the next chapter we review the theory of ICA and the techniques which aid it in separating variables (signals) from data.

1.3 Research Aim

In this thesis we focus on the application of Independent Component Analysis (ICA). Before the aim can be stated we define the terms used more specifically. We use the term *climate process* to refer to a structurally cohesive atmospheric feature that describes a part of the behavioral variations in the atmospheric dynamics. Also, in ICA the patterns are known as *signals* or Independent Components.

The contributions that this research makes are outlined in the following statements. We first look at the *sensitivity* of the signals:

- **When Changing the Number of Signals**: Westra et al. (2010) shows that the results of ICA are sensitive to the number of signals estimated to be in the data. As a common way to determine the importance of patterns, is by using the portion of total variance of the data that they explain, we use it as a measure of sensitivity. Our contribution is therefore to determine how sensitive the signals are to changes in the number of signals separated from the data.
- **When Changing the Initial Estimate of the Signals**: A common way of classifying the relative importance of the signals is to rank them in terms of the proportion of the total variance of the data that they represent. Our contribution is to examine if any changes occur to the order of the signals when the initial estimate of the signals uses random values.

Then by examining the *time series* and *spatial maps* produced by signals:

- **Spatial Maps:** A set of spatial maps represents the spatial manifestation of a signal over time. Our contribution is to examine if the use of spatial maps aids in the interpretation of signals to identify processes.

Therefore the specific aim of this thesis is:

To determine the value of Independent Component Analysis in identifying climate processes.

1.4 Outline of Research

The work of this thesis is achieved through the following chapters:

- **Chapter 2 [Exploring the Foundations]**
  Goes into the details behind the theoretical aspects of the work. It covers Singular Value Decomposition, and how it can be used for performing Principal Component Analysis. Lastly it presents Independent Component Analysis.

- **Chapter 3 [The Sensitivity of Independent Components]**
  Examines the sensitivity of the signals in terms of the proportion of variance that they explain. The sensitivity is examined by changing the number of signals separated from the data, and by changing the initial estimate of the signals.

- **Chapter 4 [Making Sense of Signals]**
  Identifies the major climate processes and examines the results in terms of the spatial maps and the time series of the signals.

- **Chapter 5 [Conclusions]**
  Major findings of the research.

2 Exploring the Foundations

2.1 Introduction

In this section the framework to separate signals from the data using Independent Component Analysis (ICA) is discussed. We introduce the foundational techniques which aid in the preprocessing of the data, followed by a discussion of Independent Component Analysis. We conclude with a prototype which demonstrates an implementation of the techniques.

2.2 Singular Value Decomposition

Singular Value Decomposition (SVD) is a method used to split data into matrices which posses useful properties. One such use of these matrices is that they offer an alternative way to implement Principal Component Analysis which is used in this work as a preprocessing tool for ICA.

With SVD, we take our desired input data matrix X to have m rows and n columns. We consider the case where we have n variables, which each have m observations, and also that the relationship of $m>n$ holds. The SVD of X can be seen in equation 1.

$$X_{m\times n} = U_{m\times n} D_{n\times n} V^T_{n\times n} \tag{1}$$

The columns of the U and V matrices are both orthogonal (uncorrelated) and have unit length. These columns are known as the left and right singular vectors of their respective matrices. The D matrix contains the singular values of the decomposition within its diagonal, with each singular value corresponding to a common singular vector in both of its neighbouring matrices. The singular values are the square root of the eigenvalues from the covariance matrix of X, and represent the standard deviation of the singular vectors. In addition to this, they are arranged in decreasing order of variance for convenience, along with their corresponding singular vectors. The remainder of the D matrix however, contains only zeros. The result is a decomposition of uncorrelated vectors, which are arranged in decreasing order of the variance they explain.

2.3 Principal Component Analysis

In Principal Component Analysis (PCA), the data consist of a finite set of varying parts. PCA is able to separate the data into these parts according to the proportion of variance that they represent. Each part is more formally known as a Principal Component or PC of the data. The first PC that is separated, represents the largest proportion of variance that is present within in the data. Any subsequent components that are separated, each represent less variance than the previous component. Furthermore, each subsequent PC is orthogonal to its predecessor. When combined, the total variance of all the separated PCs equals the variance of the original data.

PCA can be implemented by using SVD. The same set of assumptions about the input data X are used, as those used in equation 1. To reduce the size of the data according to its variance, only some of the decomposed matrices are retained. By reducing the number of singular values and their subsequent singular vectors from U and V, we are able to reduce the variance of X and also the number of PCs. If only k singular values are retained, with k being less than both m and n, then the reduced form of the SVD of X can be expressed by equation 2. The V^T matrix that was decomposed using SVD is transposed and used to rotate the original data X, to create the set of k PCs found in the columns of Y. In this way the columns of Y are uncorrelated and are therefore orthogonal to each other.

$$\widetilde{X}_{m\times n} = \widetilde{U}_{m\times k}\widetilde{D}_{k\times k}\widetilde{V}^T_{k\times n} \quad (2)$$

$$Y_{m\times k} = X_{m\times n}\widetilde{V}_{n\times k} \quad (3)$$

Matrix	Names
U	Columns are Left Singular Vectors
D	Singular Matrix (Diagonal contains singular values)
V	Columns are Right Singular Vectors (in SVD)
	Columns are Principal Vectors (in PCA)
	Columns are Empirical Orthogonal Functions
Y	Columns are Principal Components
	Columns are Empirical Orthogonal Function Coefficients

Table 1: Common terms used in PCA.

A summary of the common matrix terms are presented in table 1. From equation 3 we can see that there is a reduction in the number of columns used, from n to k. If we then chose to reconstruct the original matrix X, from these k components, we would find that the variance of the reconstructed matrix is less than that of the original matrix. Therefore there is a trade-off between the variance retained by our reduced number of components, and any benefits we gain in using fewer components to represent our data. Furthermore, by reducing the variance of our data, there is an implicit assumption that we are not interested in components which have low variance (Stone, 2004, p110). This can be appropriate assumption to make in the field of climatology where PCs with large variances are dominant (von Storch and Zwiers, 1999, p294).

The Percentage of Variance Explained (PVE) by a Principal Component can be calculated from the singular values in D. As the singular values represent the standard deviations of the PCs, they can are squared to represent the variance of the PCs. The variances are then normalised to represent the proportion of variance that they explain, and are lastly multiplied by 100 to represent the PVE of the PCs. The formula for calculating the PVE of PC q ($q \epsilon k$) is shown in equation 4, using the square of the q^{th} singular value divided by the sum of the square of all the singular values.

$$PVE\,of\,PC\,q = \frac{D_q^2}{\sum D^2} \times 100 \tag{4}$$

There is no one method for determining the size of k. Some less formal methods include retaining only the first few PCs or the first few which explain a specific amount of variance of the original data (Jolliffe, 1986). However determining how many components, or the amount of variance to retain is case dependent. More formal methods do exist such as the one by Krzanowski (1982). The method examines how well each subsequent component is able to predict the data when compared to the degree with which the previous component was able to predict the data. When additional components are no longer able to significantly improve the prediction of the data, then only the existing PCs are retained. Overland and Preisendorfer (1982) use a different approach altogether for estimating the number of k PCs to retain. They argue that the size of k should be determined by all those PCs that are at least equal to the level of noise generated by a similar, but random dataset. Any components which have less variance than their corresponding component in the random dataset, are deemed to be suspect and are therefore excluded. They defined their method as the Rule-N function.

2.4 Independent Component Analysis

Figure 1 provides an example of Independent Component Analysis (ICA). The goal of ICA is to take a set of mixtures of signals and from them, separate a new set of signals that are distinct from noise.

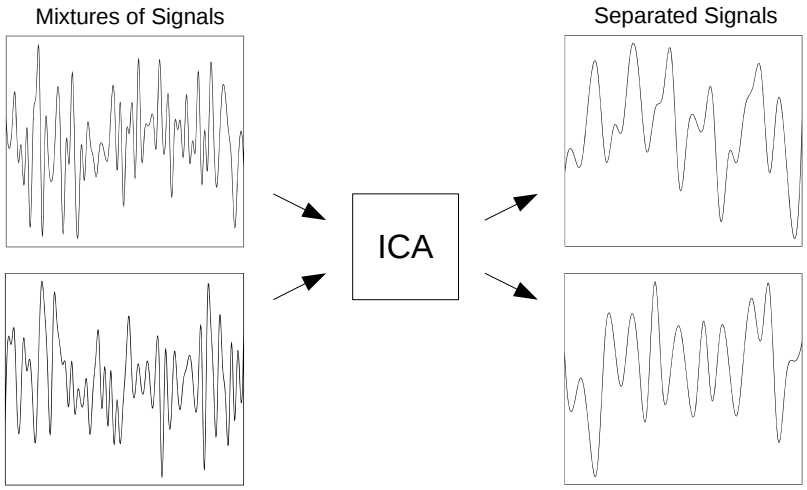

Figure 1: The function of ICA, to separate signals from mixtures of signals.

More formally ICA works by taking Gaussian mixtures of signals and from them separates a set of statistically independent non-Gaussian[2] signals. Statistical independence can be conceptually understood from figure 2. It shows that for independence to be true for two variables (A and B), they must have no information in common. As no information in the one variable can be used to predict the information in the other variable, they are therefore independent of each other. The figure also shows that when the variables share information they are therefore not independent of each other. For a more formal description of statistical independence and information the reader is referred to the work by Hyvärinen and Oja (2000).

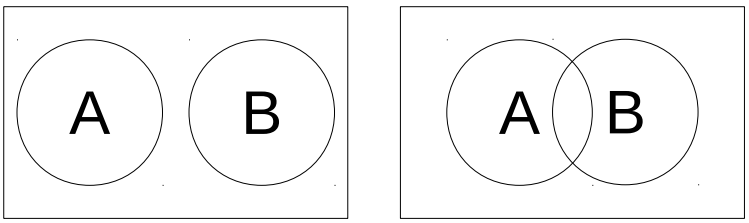

Figure 2: A conceptual diagram to demonstrate independence. [*left*] Variables A and B share no information and are therefore independent. [*right*] A and B share information and are therefore not independent.

2.4.1 The ICA Model

The more formal definition of ICA is presented by the noise-free ICA model (Hyvärinen, 1999) in equation 5, which shows how to recover the signals from the mixtures. We work backwards here, determining how the signals in S were mixed by A to form X. By determining A, ICA finds how to separate the signals from X.

$$X_{m\times n} = A_{m\times m} S_{m\times n} \tag{5}$$

Specifically X is assumed to be a matrix of the input data, which contains the mixtures of the source signals in its rows. It has m mixtures, with each mixture observed over n time intervals. The rows of S are the underlying *source signals* that were mixed within X, but can be separated from it. These source signals are also commonly known as the Independent Components (ICs) of the data. Each IC is statistically independent of every other IC, and they are therefore mutually independent of each other.

The columns of A provide us with the degree to which the source signals were mixed within X, and therefore A is known as the *mixing matrix*. Unlike the separated signals, the columns of A are not statistically independent of each other. The importance of this is that the rows of S are made independent through the columns of A.

[2]One source signal is allowed to be Gaussian, with the rest being non-Gaussian (Hyvärinen and Oja, 2000).

A closer look the at the ICA model (equation 5) shows us that it performs its separation from the signal mixtures in X. So the number of signals extracted (m) is equal to the number of rows in X. Due to the nature of ICA there are also three ambiguities associated with the ICA model that Hyvärinen and Oja (2000) point out. They state that the:

1. ***Signs of S are unknown.*** As the ICA model is composed of the A and S matrices, any sign change in one can be canceled by a sign change in the other. Therefore it is not known what sign the signals have. This can be easily remedied by multiplying the row of S and its corresponding column of A by -1 should it be found necessary to change the sign of a signal.

2. ***Variances of S are unknown.*** Similarly in the ICA model the variances of the signals cannot be determined. This is due to when either the columns of A or the rows of S are multiplied by constant, then the same operation can be undone by dividing the rows of S or columns of A by that same constant. The solution to this is to make the signals have unit variance. By dividing the signals by their standard deviation and multiplying their corresponding columns of A by the same value, the ambiguity can be solved. Note that in doing this we preserve the amount of variance of input data, into and out of the ICA model.

3. ***Order of S are unknown.*** The order of the rows in S and the columns in A can change as these matrices are both unknown. To solve this, a custom ordering method can be imposed upon the separated signals. An example of an ordering method is shown in section 3.3.1.

The noise-free property of the model means that there is no explicit error term within the model to handle any noise (Hyvärinen, 1999). This has the benefit of making the model more tractable to implement than if it did include the term. A further justification for not using an error term may be that when PCA is used as a data preprocessing tool, it retains only a subset of the original dimensions of the data according to the variance that they represent (see section 2.3). By reducing the variance of the data, the implicit assumption is made that we are not interested in components which have low variance (Stone, 2004, p110). So by using PCA as a data preprocessing tool we can make the claim that we are removing some of the low-variance noise inherent in our data prior to performing ICA.

The interpretation of the signals by the ICA model is reliant upon the type of data within the dimensions of the input data. In the ICA model described in equation 5, the dimensions of the input data are of space by time ($m \times n$). By using the input data with space by time dimensions, an underlying assumption about the form of signal independence is made. The importance of the different forms, is that one of them is always implicitly used when the ICA model is applied to data. The work by (Stone, 2004, p105) explains the two forms of independence:

- **Temporal ICA**. In the ICA model, the dimensions of the input data (X) are of space by time. When ICA is applied to the data, the signals separated are *temporally* independent of each other. This form of independence is known as temporal ICA (tICA).
- **Spatial ICA**. Using the transpose of the input data (X^T), the dimensions of the data are of time by space. On the application of ICA to the transpose of the input data, *spatially* independent signals are be separated. This form of independence is known as spatial ICA (sICA).

2.4.2 Number of signals to retain

In the area of meteorology and climatology there are also many different ways of determining the number of ICs to retain from the data. A common method, is to separate the same number of independent components as principal components. The PCs can be used as a reduced form of the data for input to ICA if the Central Limit theory is assumed. The theory states that the sum of non-Gaussian signals will tend towards a Gaussian distribution (Stone, 2004, p56). Therefore the PCs from PCA can be considered valid mixtures. One method for determining the number of PCs to retain for ICA is based on the proportion of variance that the PCs explained of the data (Fodor and Kamath (2003), Lotsch et al. (2003)), while another was to chose the number based on computational performance (Basak et al., 2004). Aires et al. (2000) choose to adopt the methodology offered by Nadal et al. (2000), which suggested retaining only a few strong PCs which would allow an adequate number of ICs to be separated.

In Scholz et al. (2004), they were interested in extracting leptokurtic signals for use in metabolic fingerprinting. Their work considered selecting the best number of ICs to use, by calculating the kurtosis of the ICs per each set of PCs. They then compiled a plot of the number of ICs with leptokurtosis against the current number of PCs being used. Each time they extracted the same number of independent components as they had PCs. The number of leptokurtic signals was at a maximum when 6PCs were used, while after 8PCs the minimum number of leptokurtic signals was seen to persist. Therefore by calculating the kurtosis for each IC per set of PCs, they were able to decide on the best number of PCs and subsequently, the best number of ICs to extract from their data.

Koch and Naito (2007) propose a combined method grafting PCA and ICA together. The method offers a trade off between lower dimensional data and the information that it contains. They propose the use of a non-parametric method which uses kurtosis and skewness to determine the number of ICs to retain. The advantage of their method is that it requires no prior information on the data to determine the number of components to retain.

2.4.3 The FastICA algorithm

The FastICA[3] algorithm by Hyvärinen and Oja (1997) provides a means for implementing the ICA model (equation 5). We use the R statistical programming language (R Development Core Team, 2010) and so use the FastICA R package (version 1.1-11)[4].

FastICA separates signals from mixtures over a number of iterations, until a predetermined threshold is reached. To accomplish this it first reduces the number of parameter to estimate in ICA by preprocessing the data using:

- **Centering.** Removes the column means of the input data, and is a requirement for the whitening process.
- **Whitening (or Sphering)**. Requires the input data be centered before it decorrelates the data, which normalises the data along its principal component axes. In doing so it reduces the number of parameters to estimate by ICA. It also provides the option to reduce the number of dimensions of the data before performing ICA. Singular Value Decomposition is one method by which the data can be whitened.

Having performed the preprocessing steps, FastICA then separates the signals from the data. It does not calculate A directly, but rather calculates the *unmixing matrix* (W) first according to equation 6. Note that the equation assumes column wise signals.

$$X_{n\times k}W_{k\times k} = S_{n\times k} \tag{6}$$

Having obtained W it then calculates A according to the inverse of W as $A = W^{-1}$.

2.5 Prototype of ICA

A prototype is analogous to a pilot study or a test case, where the feasibility of a larger more complex task can be gauged by creating a smaller replication of it. As the prototype is on a smaller scale it can be constructed in a short period of time, which allows for the rapid testing of its methodology and the software tools that are required to implement it. The early and quick testing makes the prototype ideal for identifying problems early in development.

Fodor and Kamath (2003) have produced an *Artificial Model* which they use as a prototype to gauge how successfully ICA is able to separate climate signals from a mixture of signals. We replicated their Artificial Model as our prototype, so that we could gain an understating into the details of their implementation. These included their preprocessing steps, methodology, problems, solutions, and analysis methods. These were investigated so as to provide a foundation for the research that later work could be built upon.

[3] Available online: http://www.cis.hut.fi/projects/ica/fastica/index.shtml
[4] Available online: http://cran.r-project.org/

2.5.1 Input Signals

The prototype creates a temperature and combined volcano signal, blends them into two mixtures, then separates them from the mixtures using ICA. The temperature signal was represented by the sinusoidal function in equation 7, and was simulated over a period of 264 months ($n = 264$) from January 1979 to December 2000.

$$Temperature\,signal = sin\left(\frac{1:n}{5}\right) \tag{7}$$

The combined volcano signal is constructed from the addition of two known eruptions (El Chichón and Pinatubo). Each of the volcanoes was modeled according to the equation in equation 8. The equation shows the inactivity of a volcano, followed by a cooling and warming period.

$$volcano = \begin{cases} \frac{-\triangle T_{max} t}{t_{ramp}} & t = t_{erupt}, \ldots, t_{ramp} \quad : Cooling\,period \\ -\triangle T_{max} e^{-\frac{t - t_{ramp}}{\tau}} & t = t_{ramp} + 1, \ldots, n \quad : Warming\,period \end{cases} \tag{8}$$

The t_{erupt} parameter marks the t^{th}month (out of n) in which a volcanic eruption occurs. The cooling proceeds at a rate of $\triangle T_{max}$ until it stops at t_{ramp}. At this time the exponential warming occurs with its exponential decay time represented by τ. The combined volcano signal is the sum of the effects of the El Chichón and Pinatubo eruptions, which are each modeled by the equation in equation 8, with their arguments shown in table 2. After being modeled they are added together to form the combined volcano signal.

	El Chichón	Pinatubo
n	264	264
$\triangle T_{max}$	0.32	0.72
t_{ramp}	59 (20+39)	161 (14+147)
t_{erupt}	39	147
τ	30	30

Table 2: Arguments for the parts of the combined volcanic signal. The values for t_{ramp} are adjusted from the original work to reflect the time at which the warming starts.

2.5.2 Methodology

The temperature and combined volcano signal were mixed according to the mixing matrix shown in figure 9 while using the ICA model in equation 5. The two source signals and their corresponding mixtures are displayed in figure 3.

$$A = \begin{pmatrix} 1.0 & 0.4 \\ 0.5 & 0.6 \end{pmatrix} \tag{9}$$

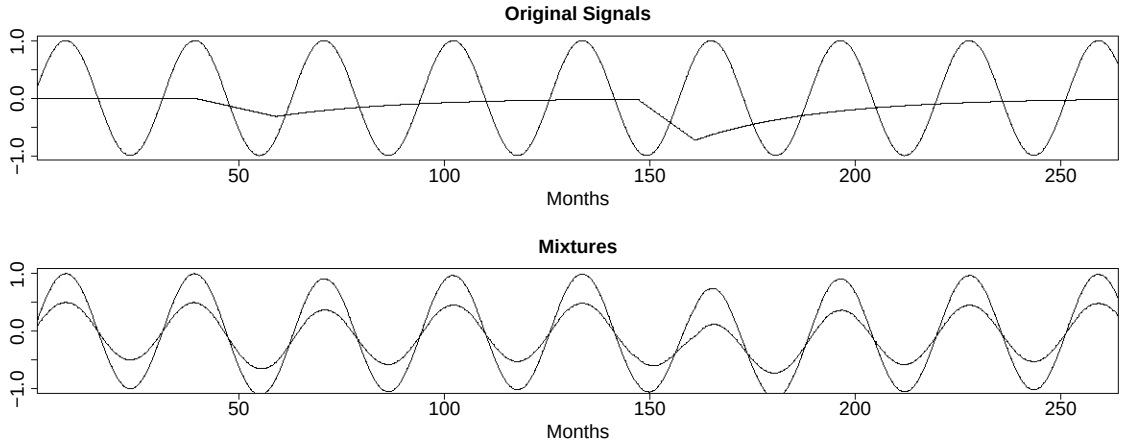

Figure 3: [*top*] The temperature and combined volcanic signal. [*bottom*] The two mixtures.

Following the mixing of the temperature and volcanic signal, ICA was applied to the mixtures to separate out the original signals. We used the FastICA algorithm (Hyvärinen and Oja, 1997)(section 2.4.3), but as we assumed row wise variables the input data and output signals were transposed.

2.5.3 Results of the Prototype

The results of the separated signals can be seen in figure 4, with the arguments required to replicate the results outlined in table 3. The extracted components are similar to those produced in the work by Fodor and Kamath (2003) within their figure 3. Note that in our work we have chosen to use ICA for the statistical independence of the signals, and therefore we use the parallel argument to the FastICA algorithm. Interestingly the separated signals are not exactly the same as those used to create the original mixtures (top image in figure 3). One reason for this may be that the source signals, while independent in theory, may not have been entirely independent in construction. The reason for this is that they may have been correlated in time during the period examined. A solution could be to use a longer time period for the analysis. This would hopefully allow the original signals to be more independent of each other, and therefore allow the separated signals to be better separated from each other as well.

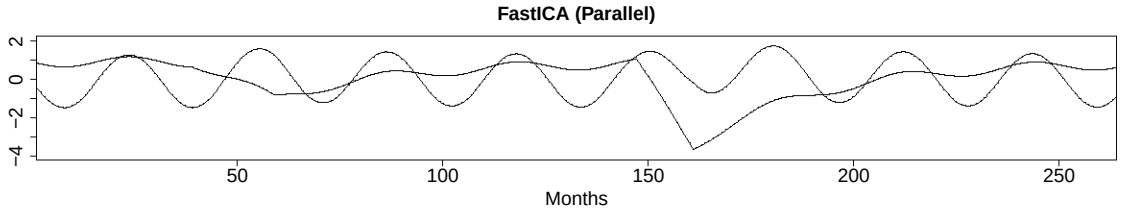

Figure 4: The signals obtained when using the parallel argument to the FastICA algorithm.

Description	**Parameter (variable)**	**Argument (value)**
Number of Signals to Extract	n.comp	2
Algorithm Type	alg.type	parallel
Contrast Function	fun	logcosh
Method	method	C
Distribution Type	alpha	2

Table 3: The arguments to FastICA needed to visually replicate the results using Independent Component Analysis. All other the parameters were assigned default values.

3 The Sensitivity of Independent Components

In the last chapter we reviewed the fundamental theory of Independent Component Analysis (ICA) and the methods which aid it in separating signals from data. In this chapter we now examine the sensitivity of the signals when we change the number of signals separated, and when we change the initial estimate of the signals. As any change in the signals can directly impact our ability to attribute them to climate processes, determining the sensitivity of the signals to these changes is important.

Specifically we first look at the effects that changing the number signals assumed to be represented by the data, has on the separated signals. This is also important as it places an upper limit on the number of possible climate processes that can be identified using the signals. Secondly we look at how the initial estimate of the signals using randomly generated numbers effects our separated signals, as substantial changes in the signals may hamper our ability to reliably attribute the signals to climate processes.

Section 3.1 explains the data and preprocessing steps that will be used in later sections. In section 3.2 we show the effects on the signals when we change the number of signals assumed to be represented by the data. Lastly in section 3.3 we examine any effects that using random values to intialise the unmixing matrix has on the separated signals.

3.1 Data and Preprocessing Steps

The data used for the application of ICA consisted of the global gridded monthly mean surface temperature data from the National Centers for Environmental Prediction (Kalnay et al., 1996). As a future application of ICA may be to evaluate global climate models, temperature data was selected due to its ability to reflect future climate change (Randall et al., 2007). The data was retained from January 1961 through to December 2000, with its mean removed and the results scaled according to each grid cell latitude to account for area coverage.

$$Weighted\,Cell\,Value = Cell\,Value \times \sqrt{cos(Cell\,Latitude)} \tag{10}$$

This produced a weighted series of climate anomalies of 144 longitudes by 73 latitudes, for a total of 480 months (40 years). The 2D spatial dimension was reduced to form a single spatial dimension, with the final data consisting of 10512 (144 x 73) spatial elements and 480 temporal elements. For brevity $X_{m\times n}$ represents the data with $m = 10512$ and $n = 480$.

Fodor and Kamath (2003) show that it is unlikely that every cell in the spatial data will represent a source signal. Therefore, before applying ICA to the data, they used PCA to reduce the dimensions of the data. PCA can also be used to remove signals with small variances which may be non-Gaussian noise. This is useful, as ICA does not distinguish between the variances of signals, and therefore it would still separate

it from the data if it was not removed before hand. So PCA was also considered in this work to reduce the number of spatial dimensions of the data. Specifically, the Rule-N function (section 2.3) was used to determine 6 out of 480 possible PCs to retain as input to ICA.

Rather than using the principal components, Stone (2004, p179-181) states that the right singular vectors can be used as input to an ICA algorithm for temporal ICA, as they already represent a whitened version of the data. In this work we chose to use this approach. Note, that the right singular vectors are also known as the Principal Vectors (PVs) of X. However, there must be a minimum of 2 singular vectors used as the mixtures for ICA as this is a fundamental constraint of the ICA model (section 2.4) . All the singular vectors selected to be retained are then used as input to the FastICA algorithm created by Hyvärinen and Oja (1997) (section 2.4.3). Note that in later sections we change number of singular vectors retained (k), but the methodology to separate them remains the same.

Step	Operation	Comments
1.	Subset NCEP Data	40yrs of Global Surface Monthly Air Temperature Data
2.	Remove 40 year mean	Produce surface climate anomalies
3.	Weight cells by latitude	Remove biases of cells with small or large latitudes
4.	$X_{144\times 73\times 480} = X_{m\times n}$	Squash 2D spatial matrix to 1D array (144 x 73 to 10512)
5.	$\tilde{X}_{m\times n} = \tilde{U}_{m\times k}\tilde{D}_{k\times k}\tilde{V}^T_{k\times n}$	Reduce to k spatial dimensions using SVD (Eg: $k = 6$)
6.	$J_{k\times n} = row\,mean(\tilde{V}^T_{k\times n})$	Retain the row-wise means of the singular vectors in ($J_{k\times n}$)
7.	$\widetilde{V}^T_{k\times n} = A_{k\times k}S_{k\times n}$	Use centered right singular vectors as FastICA mixtures
8.	S have unit variance	Remove variance model ambiguity of S (scale A accordingly)

Table 4: Summary of methodology and preprocessing steps

Although the FastICA algorithm is able to reduce the dimensions of the data, it was found that the covariance matrix of $X_{m\times n}$ it requires was too large to create with the current computational constraints. To solve this SVD was used to reduce the dimensions of the data before execution of the FastICA algorithm. In all work, the same number of independent components were separated as there were mixtures in the data, so no additional dimension reduction was performed by the algorithm. Also, any rotation of the data by the FastICA whitening step would create another mixture from which to separate the signals.

The arguments to the FastICA are given in table 5. As we have assumed that the mixtures of our signals are in the rows of our input data while the FastICA algorithm rather assumes that they are in the columns, we transpose the singular vectors before separating the signals from them.

Description	Parameter	Argument
Signal Mixture	X	$\tilde{V}_{n\times k}$
Number of Signals to Extract	n.comp	k
Algorithm Type	alg.type	parallel
Contrast Function	fun	logcosh
Method	method	C
Level of Output Information	verbose	False*
Distribution Type	alpha	2
Normalize the rows	row.norm	False*
Number of Iterations	maxit	200*
Convergence Threshold	tol	10^{-4}*
Initialize unmixing matrix	w.init	null*

Table 5: List of arguments used for the FastICA algorithm. (Default values indicated by *)

To determine the contribution of the separated signal, we reverse our methodology from steps 7 to 5 in table 4. We first work back through steps 7, 6, and 5 in equation 11, recreating $\tilde{V}^T_{k\times n}$, adding back the row means, followed by recreating an approximation of $X_{m\times n}$. It is this approximate representation of the data that is used for calculating the variance of the signal. Note that for signal q, the q^{th} column of A and q^{th} row of S are used where $q\epsilon(1, 2, ..., k)$. Using this information the percentage of variance is calculated. The reconstructed variance of the signal is represented as a percentage of the original variance to give us the Percentage of Variance Explained (PVE) by signal q in equation 12.

$$Variance\,of\,signal\,q := Var(\widetilde{U}_{m\times k}\widetilde{D}_{k\times k}((A^q_{k\times 1}S^q_{1\times n}) + J_{k\times n})) \tag{11}$$

$$PVE\,by\,signal\,q := \frac{Variance\,of\,signal\,q}{Variance\,of\,(X_{m\times n})} * 100 \tag{12}$$

As the signals are separated from the mixtures, the cumulative PVE by the signals is equal to the cumulative PVE by the PVs, thereby placing an upper bound on the amount of variance that a set of signals can explain. Additionally, the distribution of the variance amongst the separated signals can vary, and is not constrained to equal the PVE by any of the PVs.

3.2 When Changing the Number of Signals

To use Independent Component Analysis (ICA), the number of signals must be determined before they can be separated from the data. Possible techniques for determining the number of signals to separate from the data are discussed in section 2.4.2. Here we are interested in examining what effects selecting a different number of signals from the data has on the Percentage of Variance Explained (PVE) by the signals. In particular we select the first 50 out of a possible 480 Principal Vectors (PVs) to use as input to our algorithm. The number of PVs to use is deemed to be sufficiently large to show any changes in the sensitivity of the signals. In doing so we wish to see how sensitive the PVE by the signals is to a change in the number of signals separated from the data.

We used the data specified in section 3.1, according to the steps in table 4 with the FastICA arguments as listed table 5. For the number of signals to use we selected the full range of 1 to 50 signals, at each time separating the same number of signals as PVs. An exception to separating the same number of signals as PVs, is that of separating 1 signal from the data. In this case the minimum number of 2 mixtures is enforced, and while there is one signal separated, there are two mixtures (PVs) used. So the mean of the mixtures taken at step 6 (table 4) cannot be added to the results, as $J_{2\times n}$ does not have the same dimensions as the resultant AS which are of $1 \times n$ dimensions. Therefore the PVE of 1 signal from the data does not include the mean of the mixture, and its PVE is less than that of the mixtures used.

The results of the analysis are presented in figure 5. The figure shows the how the PVE by the signals changes with an increase in the number of signals separated. The first ranked signal is substantially larger for most of the sets. Additionally the behavior of the first signal can be seen as to exist in two phases. The first phase extends from set 2 to set 28. In this phase the first ranked signal generally decreases in its PVE as the size of the sets increases.

In the second phase from set 29 to set 50, the first ranked signal appears to change erratically in PVE. Towards the last few sets it reaches a similar PVE value as the remainder of the signals in the sets. This may be due to certain PVs introducing new information which increases the PVE by the first signal, followed by the splitting of the signal which is then represented by multiple signals in the next set, but with lower PVE values.

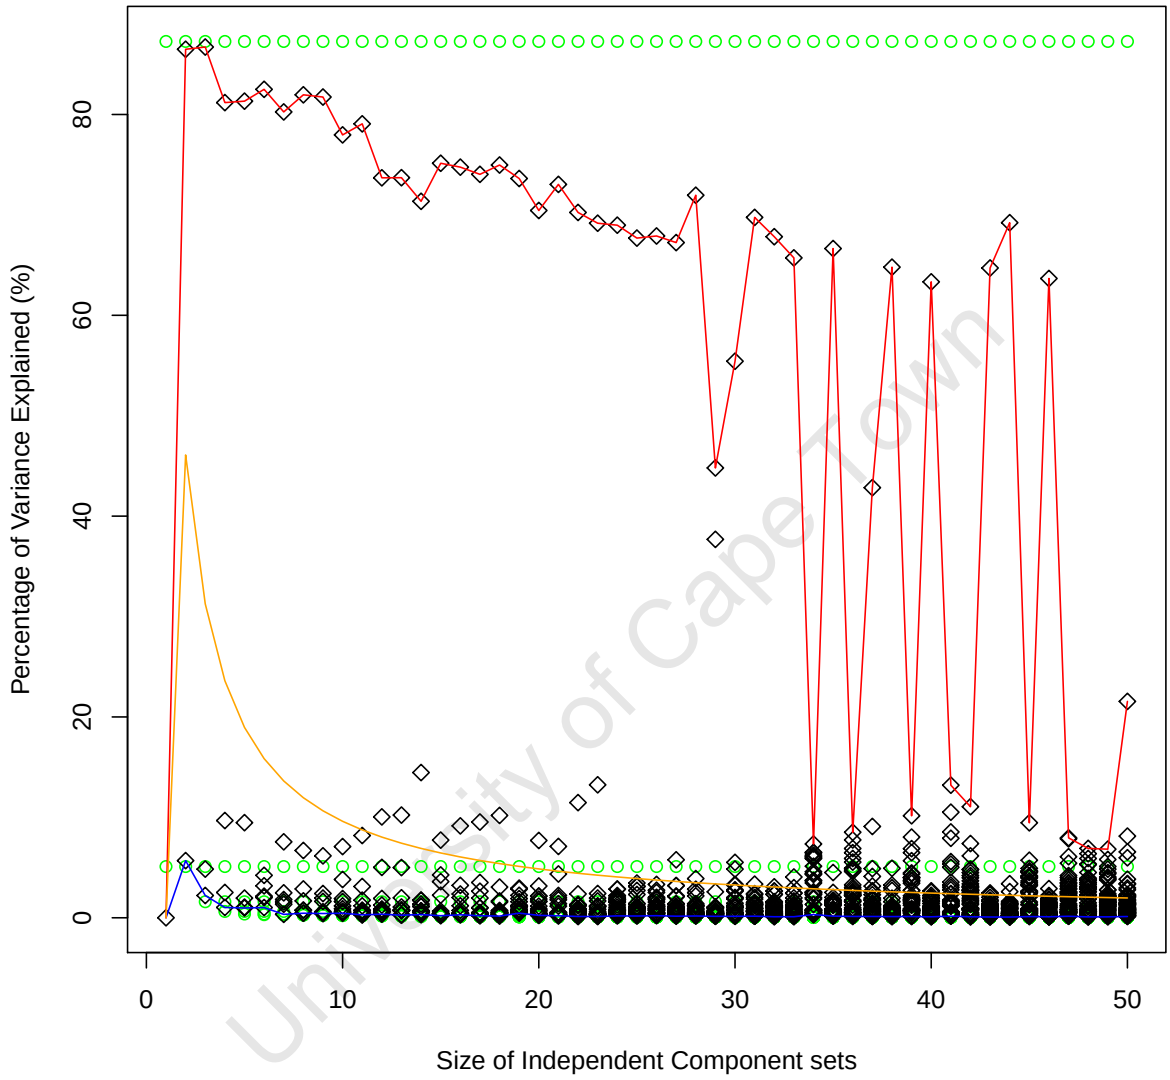

Figure 5: The changes in the variances explained by signals (diamonds) and principal components (circles) when assuming different numbers of signals in the data. The maximum, running average, and minimum variances explained by the signals are plotted as lines.

To better explain what information the PVs introduce into the mixtures that changes the PVE by the signals, we construct static maps. A static map provides an image which represents the spatial manifestation of the signal, and which spans the entire time period represented by the data. By calculating the difference between two static maps, we can see what spatial information has changed between any two signals. The equation for a static map is shown in equation 13 (Stone, 2004, p181).

$$Static\,Map\,for\,signal\,q:\ L_{m\times 1} = \widetilde{U}_{m\times k}\widetilde{D}_{k\times k}A^{q}_{k\times 1} \tag{13}$$

The q^{th} column of A is used to generate the static map for the q^{th} signal. The static map is then changed from an 1D array into a 2D image, reversing the effect of step 4 in table 4. Additionally, we undo the weighting effect (step 3 in table 4). Due to the map projection the polar regions may appear to be more important in area than they actually are. The static maps provide a way for determining the cause behind the change in PVE by signals, when the number of signals separated from the data changes.

As the first ranked signal explains substantially larger variance than most of the remaining signals in a set, we investigated it further so as to determine what information has been gained or lost by use of the additional PVs. The difference between static maps of the first ranked signal from set 28 and set 2 is shown in figure 6. We chose these sets as they are the endpoints of the first phase of the first ranked signal and have a large difference in PVE of 12.67%. The static map with the unscaled results shows that the major differences can be seen over the high latitude regions and North Africa. The reason for this may be due to seasonal cycle information being added, as higher latitude regions undergo larger seasonal changes than lower latitudes.

The static maps provide the spatial manifestation of signals over the entire time period of the data. The difference between the static maps for the first ranked component from different sets, highlights the sensitivity of the variance explained by the signal on the number of signals separated in figure 5. Therefore it is crucial to determine the number of signals to separate, as their PVE and spatial representation can change when using a different number of mixtures. The use of static maps further highlights the changes in information gained or lost when changing the number of signals separated. Lastly the variance explained by the signals can change when using a different number of signals, unlike Principal Component Analysis where the principal vectors remain at a fixed variance regardless of the number of signals separated from the data. This further reinforces the need to fully investigate how many signals to use in Independent Component Analysis, as the results of the variance explained and static maps are dependent upon the number of signals separated.

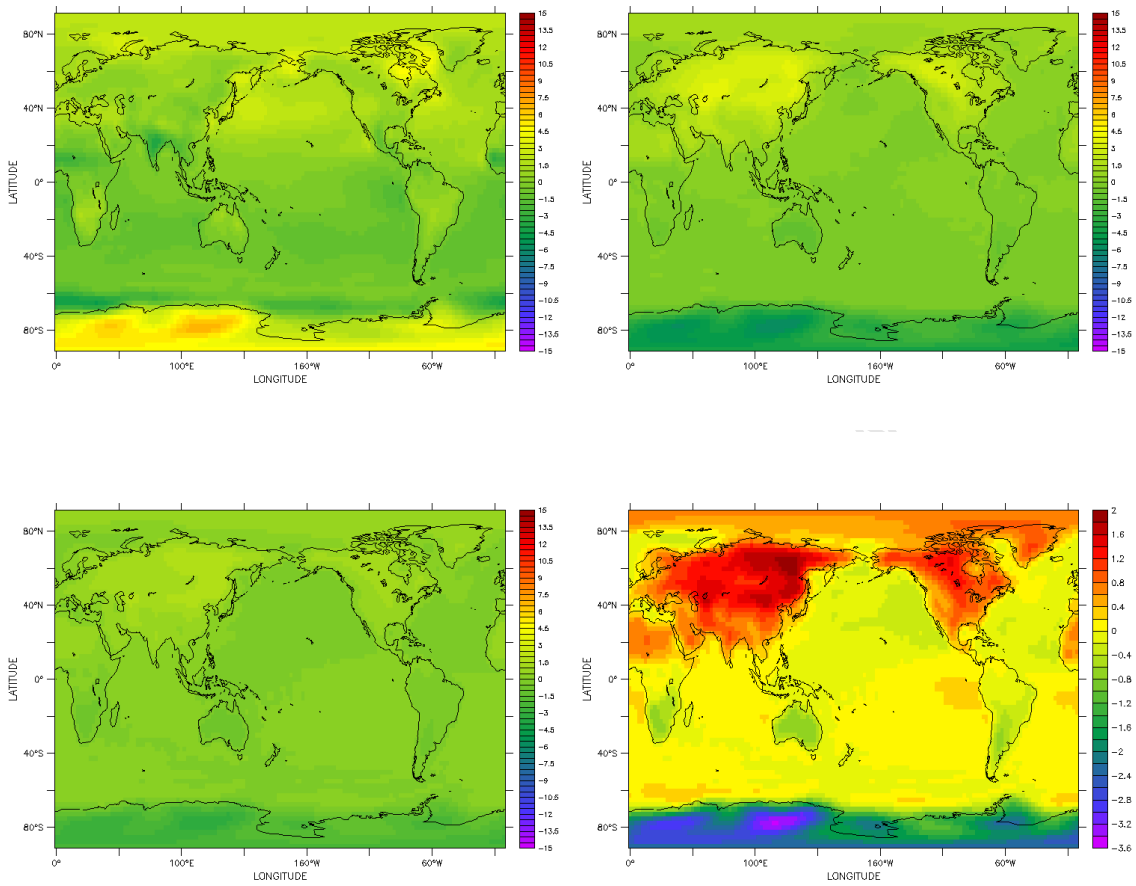

Figure 6: The static map for the first ranked signal from [*top left*] set 2 and from [*top right*] set 28. The difference between the static map of 28 and 2 is shown [*bottom left*] as scaled and [*bottom right*] unscaled. Note that as the signs of the static maps are arbitrary they are made the same for plotting purposes, and that white regions represent data that is outside of the scale.

3.3 When Changing the Initial Estimate of the Signals

We now investigate the effects caused by the initialisation of the unmixing matrix with random values on the separated signals. The unmixing matrix of the FastICA algorithm (Hyvärinen and Oja, 1997)(section 2.4.3) is initialised using a set of random values, which form the initial estimate of the unmixing matrix. This initialisation process is performed each time the algorithm is executed. As the unmixing matrix is used to separate the signals from the mixtures, it determines the structure of the signals.

The challenge is to deduce if the signals are sensitive to the initialisation of the unmixing matrix with random values. To aid in this task we employ a classification scheme which we then use to classify any structural changes in the signals. We then measure any classified structural changes using a metric known as the *mapping stability,* which acts as a measure of how sensitive the signals are to the initialisation of the unmixing matrix.

Section 3.3.1 defines the structural scenarios and the mapping stability metric. Section 3.3.2 performs a test to confirm that if any changes in the structure of the signals are observed, then they can only be due to the initialisation process. The setup of the experiments is covered in section 3.3.3. Sections 3.3.4 and 3.3.5 discuss the results of the initialisation of the unmixing matrix for sets of 6 and 10 signals respectively. Lastly, section 3.3.6 summaries the results.

3.3.1 Measuring Structural Changes

In order to measure the magnitude of the initialisation we employ a series of metrics. These metrics are used to analyze sets of signals, each from different executions of the FastICA algorithm. By comparing signals between sets from different executions, we are able to observe any differences in their structures:

- **First Maximum Absolute Correlation (FMAC).** This provides us with the magnitude of structural similarity between signals from different sets. The constraint of using the absolute correlation when comparing signals between sets is to ensure that the arbitrary sign of extracted ICs does not compromise the sign of the correlations between two signals, as identical signals can been separated from the data with opposite signs, which would result in opposite signed correlations. The constraint of using the first maximum match, is to ensure that only one signal is found to correlate the highest with the original signal. In doing so it attempts to match one signal from the current set to its equivalent in the next set.
- **Percentage of Variance Explained (PVE).** As it is derived from the signal (equations 11 and 12), a change in the structure of a signal will result in a change in the PVE by the signal as well. It therefore serves as a measure of structural changes on the signal.
- **PVE Rank.** Allows us to impose an ordering on the signals. We arrange them in decreasing order of PVE values, with a PVE rank of 1 being obtained by the signal that has the highest PVE value within a specific set. We assume that if the initialisation has no structural affect on the signals, then the ordering of the signals should be preserved throughout the different sets. The PVE rank ordering method was chosen due to the wide use of PVE in the literature with ICA and PCA (Aires et al., 2000, Lotsch et al., 2003, von Storch and Navarra, 1999, p242). Other methods considered for ordering signals were by uncertainty (Westad and Kermit, 2003), mutual information (Back and Trappenberg, 1999) and a non-Gaussian measure (Hyvärinen, 1999).

Structural Scenarios.

To simplify the analysis task using the combinations of FMAC values, PVE values, and PVE ranks, they are grouped together into structural scenarios.

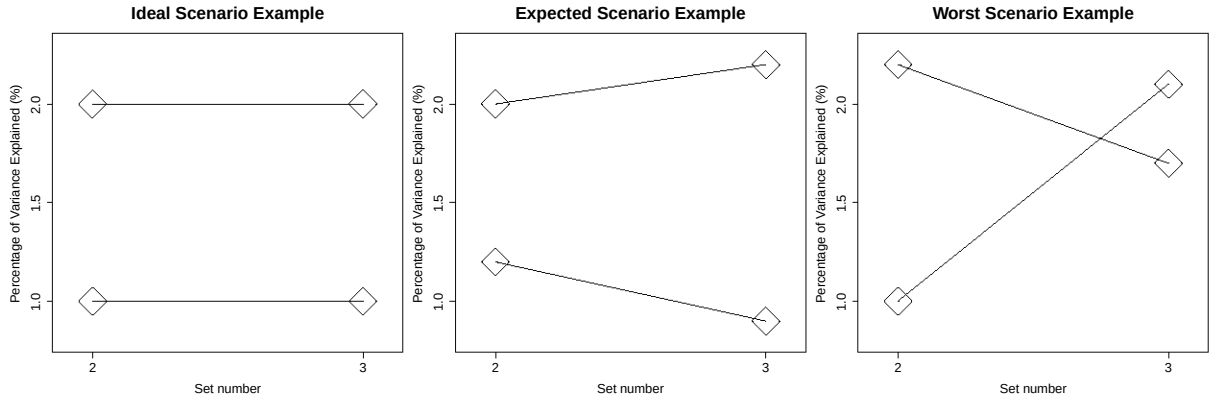

A B C

Figure 7: Examples of the different structural scenarios, comparing two signals from two sets. Links between signals are established by FMAC values, and are indicated by lines. For simplicity, the links are shown to be symmetric between sets. [A] Ideal Scenario. [B] Expected Scenario. [C] Wost Scenario.

The *ideal scenario* is if the random initialisation of the unmixing matrix has no effect on the signals, then the matching signals in two different sets should have the same PVE rank and PVE value (figure 7a). Identical PVE values between matching signals from different sets would show that the application of the algorithm to the data, does not have any effect on the signal structures or PVE values. This scenario serves as a benchmark to test other scenarios against, for we know from equation 12 that a change in a signal will result in a change in its PVE value.

The *expected scenario* is if the initialisation of the unmixing matrix causes a slight change in the signal structure upon every execution (figure 7b). The result of the slight change would be that the highest match between two signals of different sets would indicate that they are essentially the same signal, and that therefore they should have the same PVE rank. As the signals would have slightly different structures, we would not expect their PVE values to be exactly the same and they would have a range of values rather than an exact PVE value in all sets.

The *worst case scenario* is if the random initialisation effect is great enough to change the PVE ranks of the maximally matching signals from two different sets (figure 7c). The change of PVE ranks would indicate that structurally different signals could overlap in PVE values. A summary of the structural scenarios is provided in table 6.

Scenario	Match	PVE values	Matched signals from different sets have:
Ideal	Perfect	Are exact	Same PVE rank and PVE value
Expected	Maximum	Range	Same PVE rank but different PVE values
Worst	Maximum	Overlap	Different PVE ranks

Table 6: The assumptions about two matching signals from different sets with a one to one mapping.

Mapping Stability

The mapping stability metric was implemented to measure the different structural scenarios against the ideal scenario. The metric compares signals from two different executions according to their FMAC values. If two signals from different sets have the same PVE rank and match to each other by FMAC value, then they are deemed to be stable for those two executions. The *mapping stability* is therefore the count of the number of stable links over a number of executions. The total number of mappings possible is equal to the total number of executions - 1, this is due to the number of possible mappings being equal to the number of links between signals from different sets. The final metric value represents the number of counted links with the same PVE rank, as a fraction of the total number of mappings possible, assuming the ideal scenario. The result is represented as a percentage, with a high percentage indicating a close to ideal signal.

Note that although the number of mappings is compared to the ideal scenario, we expect changes in PVE to occur, so any measured PVE changes are not counted against the mapping stability of a signal. So it does not penalise a change in PVE that occurs without a change in PVE rank as well. The mapping stability therefore measures how constant the mappings remain between signals over all executions of the algorithm.

3.3.2 Validation of Hypothesis

In order show that any results we obtain are caused solely by the initialisation of the unmixing matrix using random values, we propose an hypothesis to be tested. We propose our hypothesis as:

Hypothesis: *As S is separated using W, no change in W should result in no change in S either.*

To prove the hypothesis, we conducted two tests using a fixed W matrix. The goal of each test was to hold the unmixing matrix as a constant throughout a number of executions of the FastICA algorithm, and monitor any structural changes using the metrics as defined in section 3.3.1. The two outcomes are:

- If no structural changes are observed then the hypothesis will be proven true, and the signals will be of the ideal scenario type.
- If however structural changes are observed, then the hypothesis will be proven false and the signals will be of either the expected or worse case scenario type.

The tests consisted of running 1000 executions of the FastICA algorithm on the NCEP data detailed in section 3.1. The first test separated 6 signals ($k = 6$), with the second experiment separating 10 signals ($k = 10$) both according to the steps in table 4 and FastICA arguments in table 5. Of special note is that for these experiments we used a pre-generated unmixing matrix as opposed to the default argument of null. The unmixing matrix was initialised with a set of random values, but held constant throughout all executions of the algorithm. Thereby using the same unmixing matrix for each of the executions. The specific sets size of 6 was determined by the Rule-N function and the set size of 10 was chosen to represent similar works (Eg: Aires et al. (2000), Basak et al. (2004)). The use of 1000 executions was deemed a sufficient number to allow any other affects to manifest themselves.

The mapping stability results for the first test separating 6 signals from the data showed a 100% mapping stability for all 6 of the components. From these mapping stability results we can tell that the signals belong to at least the expected scenario. As the mapping stability does not indicate that no PVE changes occur, we cannot confirm that they are from the ideal scenario with the mapping stability results alone.

To do prove they are from the ideal scenario set, we show the PVE values for the signals in figure 8. In this figure we can see that the PVE values for all the signals are constrained to be specific values, rather than ranges of values. The figure also shows the first 10 executions of the FastICA algorithm with the corresponding 10 sets of signals, which allows us to see some evidence for the ideal scenario.

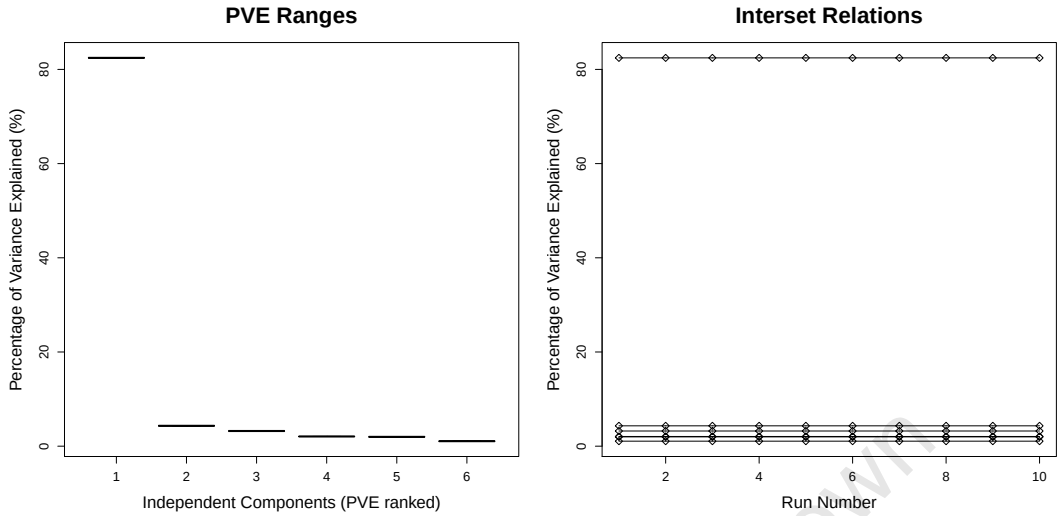

Figure 8: [*left*] The boxplot of the PVE values of the signals assuming the expected scenario. [*right*] The PVE values for the first 10 executions, with links between structurally similar signals.

This therefore confirms that the signals are of the ideal scenario, as they have the same PVE ranks (100% mapping stability) and PVE values (figure 8) throughout the executions. The results for the second test were identical to the first (not shown), with the 10 signals also found to be of the ideal scenario type. This also therefore confirms our hypothesis, as no structural changes between identical executions of the FastICA algorithm were found while maintaining a fixed unmixing matrix.

3.3.3 Implementation of Experiments

Two experiments were performed, separating sets of 6 ($k = 6$) and 10 ($k = 10$) signals from the NCEP with the data described in section 3.1, steps in table 4, and the FastICA algorithm arguments in table 5. The signal sets were separated from the data 1000 times, however unlike the hypothesis tests, the unmixing matrix was re-initialised with random values within each execution of the algorithm. The choice of using these set sizes is the same as was used in the hypothesis. The use of 1000 executions was deemed a sufficient number to allow any structural variations to manifest themselves.

Each experiment used a different 1000 executions of the FastICA algorithm. For each execution, each set of signals were separated and ranked in decreasing order of their PVE. Following this, the stability metric was calculated for all the signals.

3.3.4 Results using 6 Independent Components

The first experiment was conducted by separating 6 signals from the data as described in section 3.3.3. The stability of the components were calculated on the full set of 1000 executions, and the results are shown in figure 9.

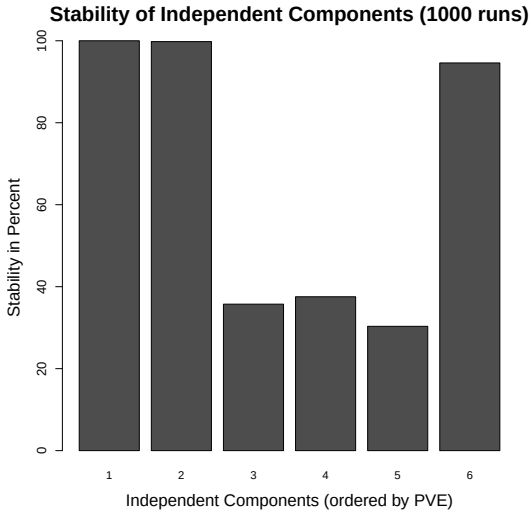

Signal	Stability (%)
1	100.00 %
2	99.80 %
3	35.74 %
4	37.54 %
5	30.33 %
6	94.59 %

Figure 9: [*left*] Barplot of stability values. [*right*] Stability Values for the 6 ICs

The first signal always maps to the first signal in the next set, thus remaining perfectly stable throughout all the executions. Interestingly though, the last component is the third most stable, while intermediate signals show less stability.

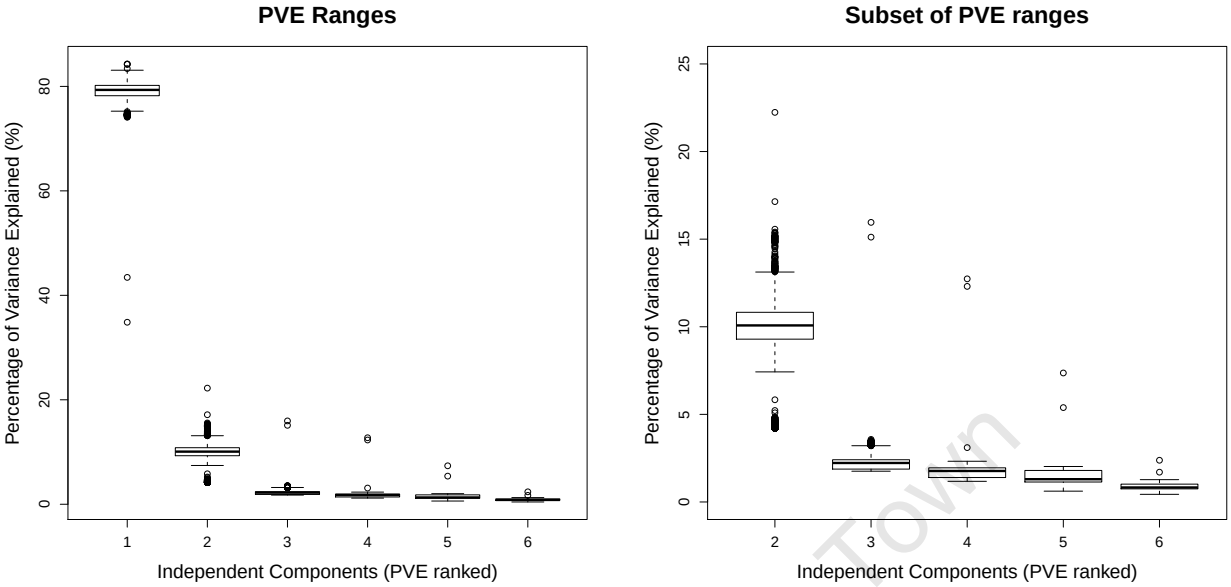

Figure 10: The PVE ranges of the signals assuming the Expected Scenario for all 1000 executions. [*left*] Shows all 6 signals. [*right*] Shows just the last 5 out of 6 signals.

Using the expected scenario for all the executions, the PVE ranges were calculated for all the signals of the same PVE rank, with the results shown in figure 10. Even with the expected scenario assumed, there is still an overlapping of PVE ranges between the signals. For example signals 4 and 5 overlap in their PVE ranges. The change in PVE values responsible for the spread in the PVE range of a signal, is not a concern, as we suspected that structural changes would result in changes in PVE values. However when the PVE ranges overlap between signals, it indicates that we can no longer associate just one signal with a given PVE range.

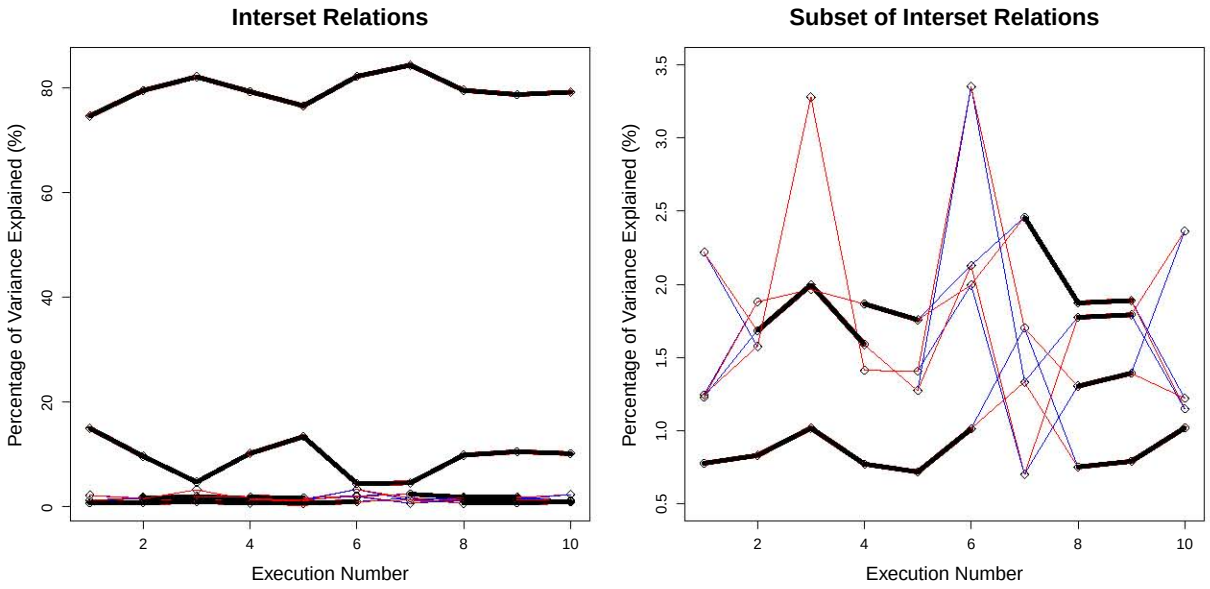

Figure 11: A visual representation of how the mapping stability metric works. The links shown, are created based on the FMAC values between signals of sets from different executions. As the metric is symmetric, *blue* lines show links between signals from the current to the next set, while *red* lines show links from the next set to the current set. If the link is common between two equally ranked signals, then the line is shown in *black*. [*left*] The PVE values for the first 10 executions. [*right*] The last 4 of 6 components for the first 10 executions, used to better show the interset relations.

From figure 11 we can see that the first IC, as ranked by PVE, is substantially larger in PVE than any subsequent signal. Also although it varies in PVE, it does not change its PVE rank in the first 10 executions. Its unchanging rank makes it an example of the expected scenario. The remaining signals with smaller PVE values are also displayed in greater detail within figure 11. The figure shows how the last 4 signals change in both their PVE values and sometimes in their PVE ranks throughout the 10 executions. An example of a signal changing its PVE rank between executions, can be seen with signal 4 in execution 4, which links to signal 5 in execution 5. This shows that the signals are not always stable in their PVE rank despite maintaining a similar structure.

3.3.5 Results using 10 Independent Components

The second experiment was implemented by separating 10 signals from 10 PVs as described in section 3.3.3, with the results shown in figure 12.

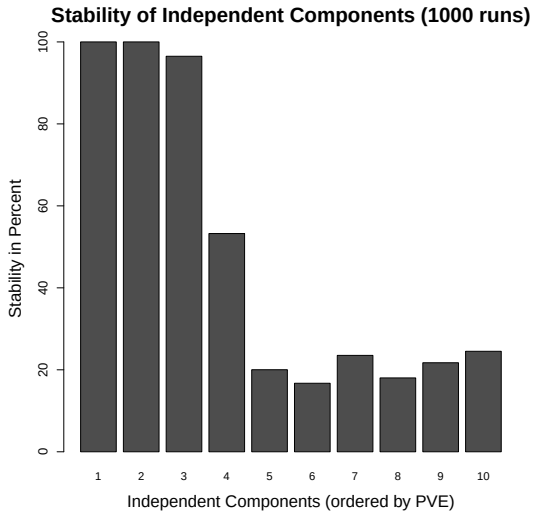

Signal	Stability (%)
1	100.00 %
2	100.00 %
3	96.50 %
4	53.25 %
5	20.02 %
6	16.72 %
7	23.52 %
8	18.02 %
9	21.72 %
10	24.52 %

Figure 12: [*left*] Barplot of stability values. [*right*] Stability Values for the 10 ICs.

The first and second ranked signals are stable throughout all the executions. However, similar to the results when using 6 signals, the remaining signals have varying stabilities. There also appears to be no relationship between the remaining signals and the PVE that they explain. In figure 13 the PVE ranges of the signals have been calculated, and a subset of the signals is shown for convenience. In the figure we can see that, except for signal 1, the other signals overlap in their PVE ranges.

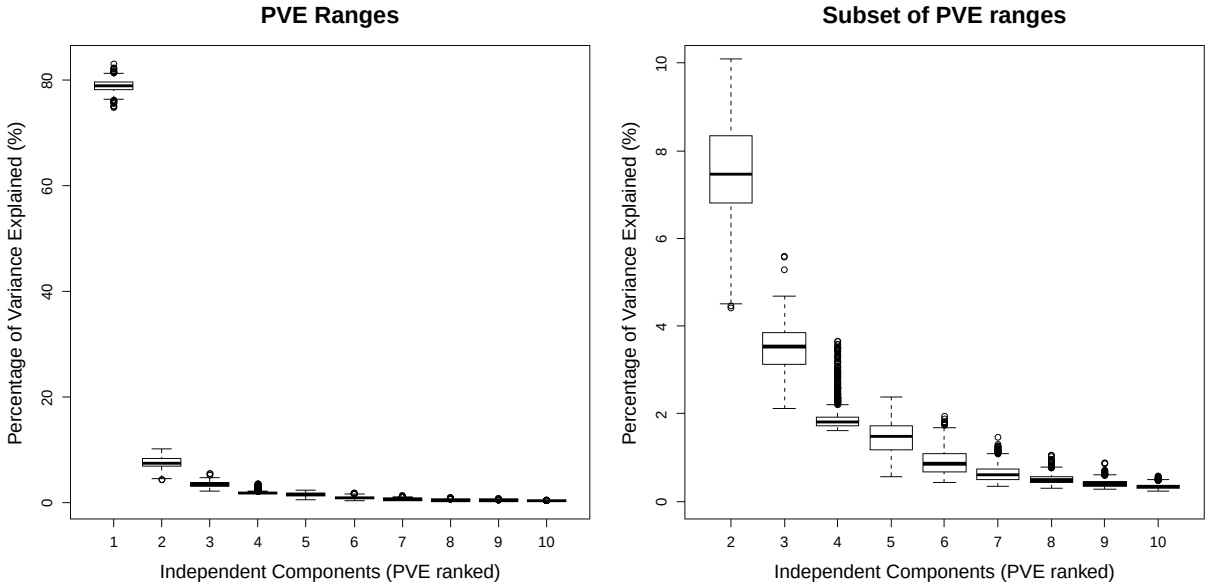

Figure 13: The PVE ranges of the signals assuming the Expected Scenario for all 1000 executions. [*left*] full set of 10 signals. [*right*] Subset of 9 signals.

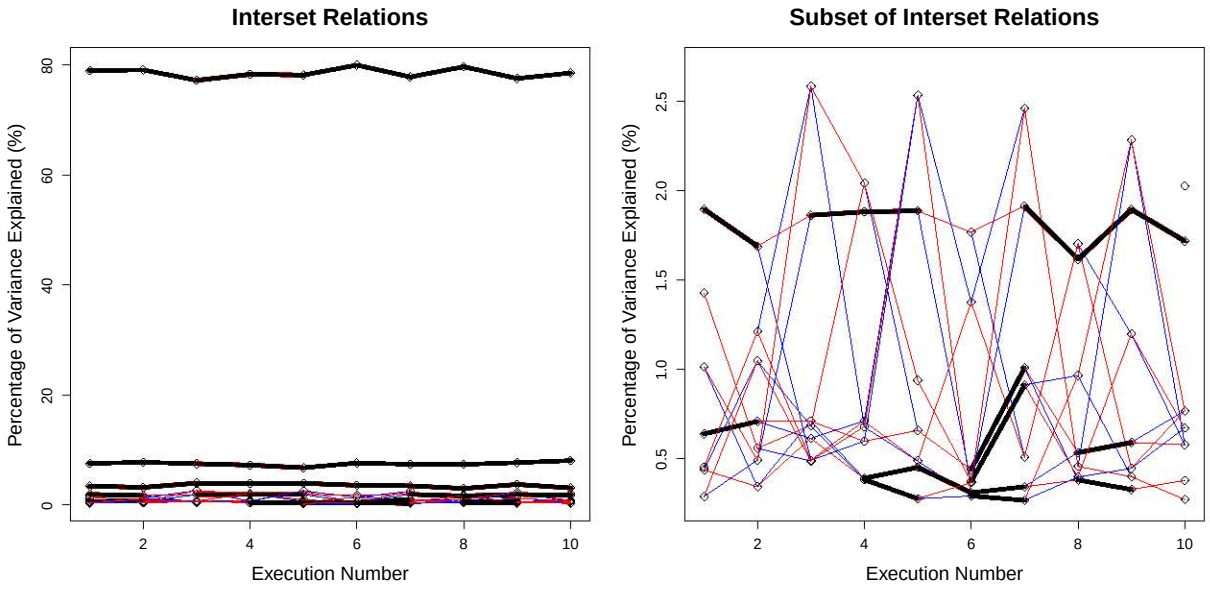

Figure 14: Similar to figure 11, but using 10 signals per execution. [*left*] The PVE values for the first 10 executions, with links between structurally similar signals. [*right*] The last 7 of 10 components for the first 10 executions.

In figure 14, the first ranked component has substantially larger PVE values when compared to the remainder of signals. The first and second ranked signal are both perfectly stable, and therefore have the expected structural scenario. The last 8 signals all show the worst scenario to some degree, with supporting evidence for the last 7 being shown in the figure for the first ten executions. These results are similar to those in the experiment using 6 signals (section 3.3.4).

The mapping stability results show that the signals are in general, more unstable than results when using 6 signals. One reason for this may be, that we overestimated the number of signals present within the data. Therefore, the FastICA algorithm maybe unable to separate the full 10 signals from the data adequately, resulting in lower mapping stabilities.

3.3.6 Summary

We investigated how robust our separation of signals from the data was, when the FastICA algorithm starts by using a random estimate of the signals. To determine that only the random estimate was responsible for any changes, a hypothesis was created and tested. We found that no other piece of code could be responsible for any effects on the signals. To quantify the magnitude of any changes on individual signals, a metric termed the *mapping stability* was developed. A signal with a high mapping stability was deemed to be structurally robust over a number of executions of the FastICA algorithm.

Two different sets consisting of a different number of signals were used, to eliminate any biases or exceptions that might have been present in one set. Each experiment used a different set, and extracted that number of signals from the data numerous times. The mapping stability for each signal was then calculated. As part of the mapping stability, the separated signals are ranked according to how much variance of the data they explain, with the first ranked signal explaining the most variance. We found that for both experiments using 6 and 10 signals per set:

- The first ranked signal explained a substantially larger variance than any subsequent signal.
- The first ranked signal was always insensitive to a change in its initialisation.
- Intermediate ranked signals generally had varying stabilities.

Although the second ranked signal in the second experiment was stable, the second ranked signal from the first experiment was slightly less than perfectly stable. This suggests that the second signal may be not always be stable and therefore a larger number of executions may reveal this in the second experiment.

From these results, there does not appear to be any fixed relationship between the rank of a signal and its stability, so explaining more variance does not necessarily imply that a signal will have a greater stability. Two additional experiments were conducted to determine any effects other parameters could have had on the results. These were done by modifying the number of iterations before convergence to independence, and then by modifying the alpha value. No substantial change in the results was observed.

The change in variance explained by the signals is not a concern. We expected that changes in the structure of the signals, due to the random initial estimate, would cause changes in the amount of variance explained by the signals. However when the variance ranges of two signals overlap, then the ordering of them may no longer be consistent between identical executions of the FastICA algorithm. The implication for the research is, that ordering the signals by the amount of variance that they explain is not reliable for any signal other that the first ranked signal. So for example if we rank components by variance, we cannot say that the same 3^{rd} ranked signal will always explain the 3^{rd} largest amount of variance. This is because its rank may change with the next execution of the algorithm.

The first limitation of the experiments is tied to the way in which the variance of the signals is calculated. The calculation is based on the assumption that the same number of signals as mixtures is always used. If less signals that mixtures are separated, the mean for the mixtures cannot be added back to the results. This results in a loss of variance, and the implication is that the variance explained by the signals can no longer be calculated precisely. Secondly, the stability results may be platform or implementation specific, with change in one of these possibly resulting in different mapping stabilities for the signals. Lastly we may be overestimating the number of non-Gaussian signals to separate from the data, and therefore the results are understandably unstable.

4 Making Sense of Signals

4.1 Introduction

Having outlined the Independent Component Analysis model and the sensitivity of its results, we now look at interpreting the signals in terms of the climate processes that they represent. The literature can be broadly categorized into the following areas:

- *Comparison of signals to previously defined climate indices (eg: Niño 3.4)*
- *Viewing of the spatial manifestation of the signals*

Comparison of signals to previously defined climate indices
While comparing the spatial patterns of the signals to the work in the literature is a means of matching the signals to climate processes, the comparison of signals to climate indices offers a more quantitative way in which to measure the strength of the match. A common form of the comparison is done through the correlation of signals to previously defined climate indices. Westra et al. (2010) used this approach when analysing their data.

In contrast to correlating the signals to the climate indices is the correlation of the columns of A to the climate indices. In the research covered by Aires et al. (2000) they correlate the columns of A, termed the *temporal base functions,* against the climate indices while also allowing a time lag between them. The time lag allows additional time for the effects of a signal to manifest itself. In the case of Fodor and Kamath (2003), their use of multiple levels of temperature data and a time lag also allowed upper atmospheric changes, time before the recording of the index reflected the changes. An alternative to correlation approaches, is the use of Linear Regression by Basak et al. (2004). They linearly fitted the columns of A to the NAO index and retained the best two components for further analysis.

Viewing of the spatial manifestation of the signals
To facilitate the matching of the signals to climate processes, the spatial representation of signals can be used. These are particularly useful in that they provide a visual representation of the signal (or its corresponding column in A) in the form of an image, and so allow for the linking of patterns in the image to climate processes using the knowledge of them from the literature. They provide a static image, which represents the spatial coverage of the signal over the entire period represented by the data. Within this analysis method there are a further two subdivisions.

The first method views the signal as an image as is the case with Fodor and Kamath (2003), Basak et al. (2004), and Aires et al. (2000). To accomplish this, the signals have their dimensions changed from a 1D series to a 2D image but whether they have the area weighting operation undone is unknown. The image

sizes and signals vary between the different works reviewed, as the techniques for separating the signals also vary, using sICA, stICA, and a non-linear ICA method proposed by Bell and Sejnowski (1995).

The second method produces a signal to data correlation image, which highlights the strength of the relationship between them. This was done by both Westra et al. (2010) and Lotsch et al. (2003). Lotsch et al. (2003) in the context of NDVI data, further takes the signal correlations and using sICA investigates them using a bootstrapping technique. This has allowed them to highlight only regions in the image with significant correlations. In addition to this they review the tICA form, and create images using the columns of A as opposed to using the signals. They find that tICA removed the cyclic seasonal patterns from the non-cyclic patterns when it separated the signals from the data.

The main advantage of using the second method is that it offers a more formal method for validating any patterns visible within an image. We now draw attention to the mapping stability results (section 3.3.6), which showed that signals can differ in structure between identical executions of the FastICA algorithm. Therefore the ability to show only significant correlations may also provide a means for overcoming structural changes when representing spatial images.

Though the comparison to climate indices and the generation of spatial images aids in the analysis of the separated signals, there is still the task of visually examining the results and comparing them to the literature. This task is inevitable, as there is currently no fully automated method for separating and matching signals to climate processes. Furthermore, Westra et al. (2010) states that the existing climate processes may be derived from other identification methods, which may make the results from ICA difficult to associate to them. An example of this is the Antarctic Oscillation which is derived from the first PC (see section 4.4.5). We propose a technique to augment the existing techniques, by creating spatial manifestations of the signals through time.

Section 4.2 presents the spatial maps analysis technique. Section 4.3 covers the setup of the experiments, while section 4.4 discusses the results of the spatial maps. Section 4.5 concludes with the major findings.

4.2 Spatial Maps

For the spatial images the current techniques create static images in time. We propose an adaption of these techniques into spatial images which change over time according to the signals. We present the technique to create spatial maps in equation 14 as a post-processing step after those shown in table 4. Here the q^{th} set of n spatial maps uses the q^{th} column of A and q^{th} row of S, with the single spatial dimension ($m = 10512$) being transformed into a 2D image (144×73) for viewing.

$$Spatial\ Maps\ from\ signal\ q: \ M^q_{m\times n} = \widetilde{U}_{m\times k}\widetilde{D}_{k\times k}((A^q_{k\times 1}S^q_{1\times n}) + J_{k\times n}) \tag{14}$$

The advantage to using spatial maps is that the literature can be first searched for important events, say El Niño years. The timing of these events can then be sought within the spatial maps at the corresponding time. If there is a corresponding pattern present in the spatial map at the same occurrence as the atmospheric event, then there is the possibility that the signal represents the climate process. The link is not conclusive, but rather serves to augment existing techniques by reinforcing the evidence for the link between the signal and the climate process. The spatial maps are therefore designed to facilitate the association of signals to climate processes. Note that we have reversed the area weighting operation performed in the preprocessing steps (table 4). Also, due to the map projection the polar regions may appear to be more important in area than they actually are. As the processes vary in their PVE, the scale which they are represented by may change between signals to accurately reflect the processes.

4.3 Setup of Experiments

To demonstrate the application of the spatial maps defined in section 4.2, we create three experiments. The experiments separate a set of signals from the data as defined in section 3.1 according to the steps in table 4 with the arguments to the FastICA algorithm (Hyvärinen and Oja, 1997) (section 2.4.3) defined in table 3. For the first experiment 4 signals ($k = 4$) were separated, in the second experiment 6 signals ($k = 6$), and in the third experiment 10 ($k = 10$) signals were separated. We chose 4 signals to compare our results to the 4 signals Westra et al. (2010) uses. The use of 6 signals is due to the results from Rule-N function, while 10 signals is used to show the results from similar work (Eg: Basak et al. (2004)) using larger number of signals.

Using the results from the experiments the signals are compared to the Seasonal Cycle, North Atlantic Oscillation (NAO), Niño 3.4, and Antarctic Oscillation (AAO) climate indices. We chose to use the absolute correlation values between signals and climate indices, due to the sign ambiguity of separated signals. Following from this, the spatial maps for the signals were created.

4.4 Results of Separation

4.4.1 Correlations with Climate Indices

The correlations with the climate indices was conducted as in section 4.3, and the results for the highest absolute correlations are shown in table 7. The PVE values are calculated over the entire 40 year period of the data, and so represent the PVE of the entire signal rather than only a single month or period.

Index	Set Size (k)	Signal	AC	PVE (%)
NAO	4	4	0.29	1.03 %
	6	6	0.43	1.02 %
	10	7	0.44	0.89 %
Niño 3.4	4	4	0.05	1.03 %
	6	3	0.22	3.18 %
	10	8	0.38	0.65 %
AAO	4	3	0.08	2.52 %
	6	3	0.10	3.18 %
	10	8	0.13	0.65 %

Table 7: Best Matching signals to climate indices. (AC represents the Absolute Correlation value.)

A point of interest in the results is that the same signal in a some sets was found to best match two different climate indices. An example of this is where signal 4 from 4 was found to best match both the NAO and Niño 3.4 indices. Though the correlation value was low in the case of the AAO match, the reason for the same signal matching two indices may be due to us only using three climate indices, and therefore using more indices may have allowed the signals to better match other climate processes.

There also appears to be no substantial improvement in the absolute correlation strength when increasing the number of signals used from 6 to 10 signals. This suggests that there may be a limit at which separating more signals from the data no longer improves the correlation strengths of the signals.

In addition to these findings, we also found a step in signal 3 of 6 and signal 9 of 10. The step in these signals corresponds to the time when NCEP switched over to using full satellite data (Kalnay et al., 1996). Thus the step is most probably showing the effect of the data change. No signal in the set of 4 showed any step wise shift. Signal 3 and 9 are shown in figures 15 and 16 respectively. Both signals have an additional 12 month smoothing filter applied which shows the change in 1979.

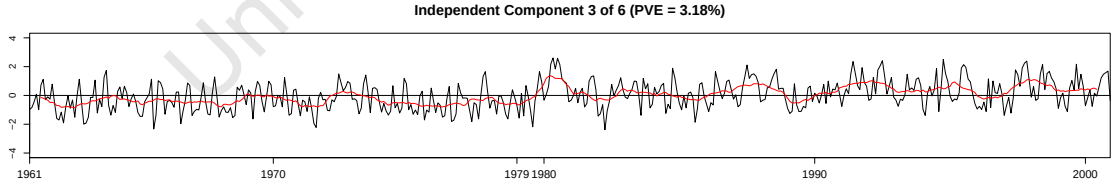

Figure 15: The third signal by PVE rank, from the second experiment separating 6 signals.

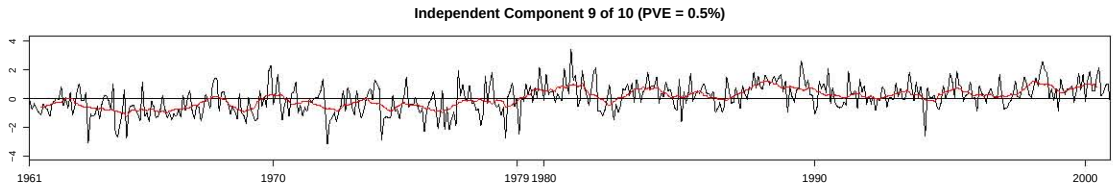

Figure 16: The ninth signal by PVE rank, from the third experiment separating 10 signals.

4.4.2 Seasonal Cycle

The Seasonal Cycle is most clearly represented by the first signal from each experiment (section 4.3) with the signal shown from experiment 2 in figure 17 and its corresponding spatial maps for the year of 1961, in figure 20. Due to the ICA model ambiguity the signals may have the inverse sign of the climate process which they represent, with the corresponding spatial maps also reflecting the inverse of the climate process.

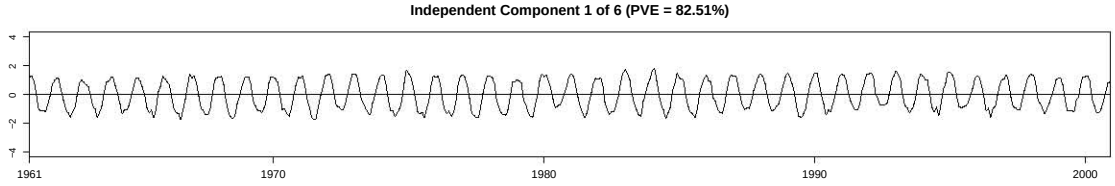

Figure 17: The first signal by PVE rank, from the second experiment separating 6 signals.

The seasonal signal is identified by the hemispheric changes in temperature throughout the course of the year. The warming in one hemisphere is synchronized with cooling in the other hemisphere. This see-saw pattern in temperature allows us to link the first component from all the experiments to the Seasonal Cycle. Further evidence that the first signals from the sets represent the Seasonal Cycle can be seen in their high PVE values (Eg: 82.51 %), as we would expect the seasonal cycle to explain a large portion of the total variance.

A mainly ocean based seasonal cycle was identified as the second signal from both the set of 4 and 6 signals. The second signal from the set of 6 is shown in figure 18. The ocean based seasonal cycle was identified in the set of 10 signals, but it appeared to be split amongst the second and third ranked signals. The largely ocean based seasonal cycles generally show a 2 month shift behind the land based seasonal cycles of their respective sets, and this shift is most likely due to the thermal inertia of the oceans (Meehl et al., 2007).

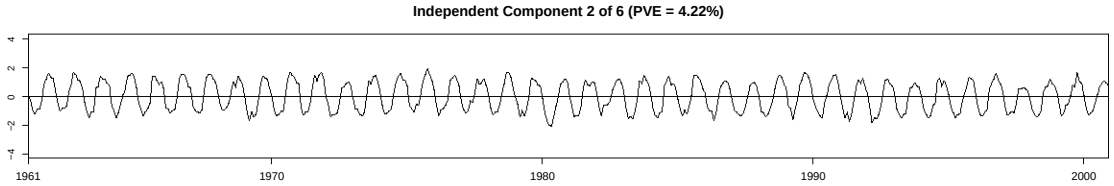

Figure 18: The second signal by PVE rank, from the second experiment using 6 signals.

In addition to these results, two 6 month cycles were found in the set of 6 signals. Both signals represent an asymmetric part of the seasonal cycle which is probably caused by snow cover. As the snow cover takes time to melt, it may be lengthening the spring warming effect, which introduces asymmetry into the seasonal cycle. The spatial maps from signal 4 indicate the effect over the mid-latitudes, while the spatial maps from signal 5 indicate the effect over the high-latitudes. We show the fourth signal from the set of 6 in figure 19 and the corresponding spatial maps for a year in figure 22. As the results of the two signals are similar the results for signal 5 are not shown.

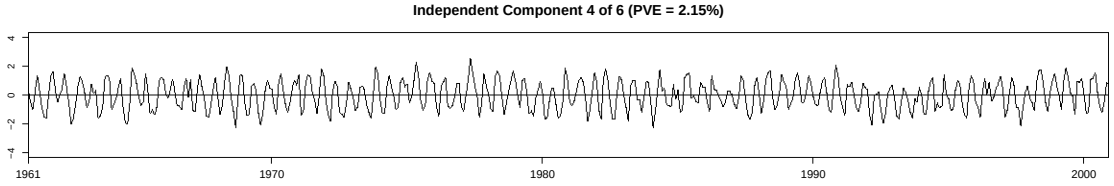

Figure 19: The fourth signal by PVE rank, from the second experiment using 6 signals.

None of the signals from the set of 4 showed the asymmetry of the seasonal cycle. While the signals from the set of 10 showed only the high-latitude effect in signal 4, with the mid-latitude effect most likely split amongst other signals.

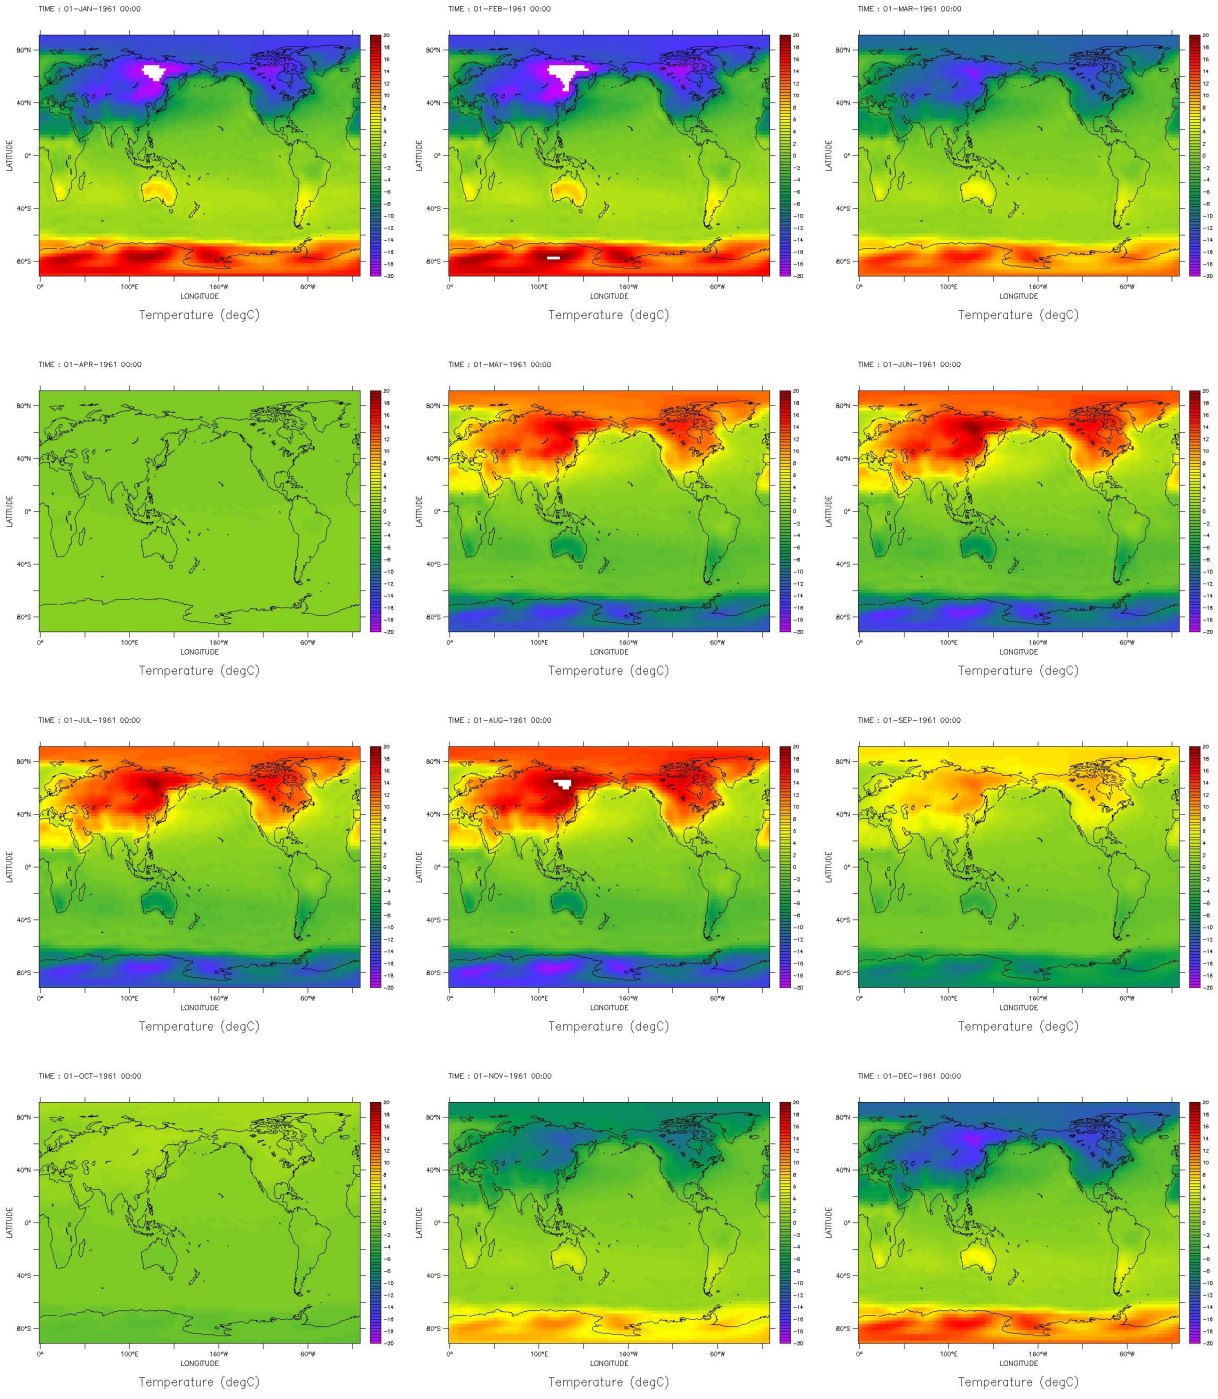

Figure 20: Spatial maps of IC 1 from the second Experiment and for the year 1961 (top left to bottom right).

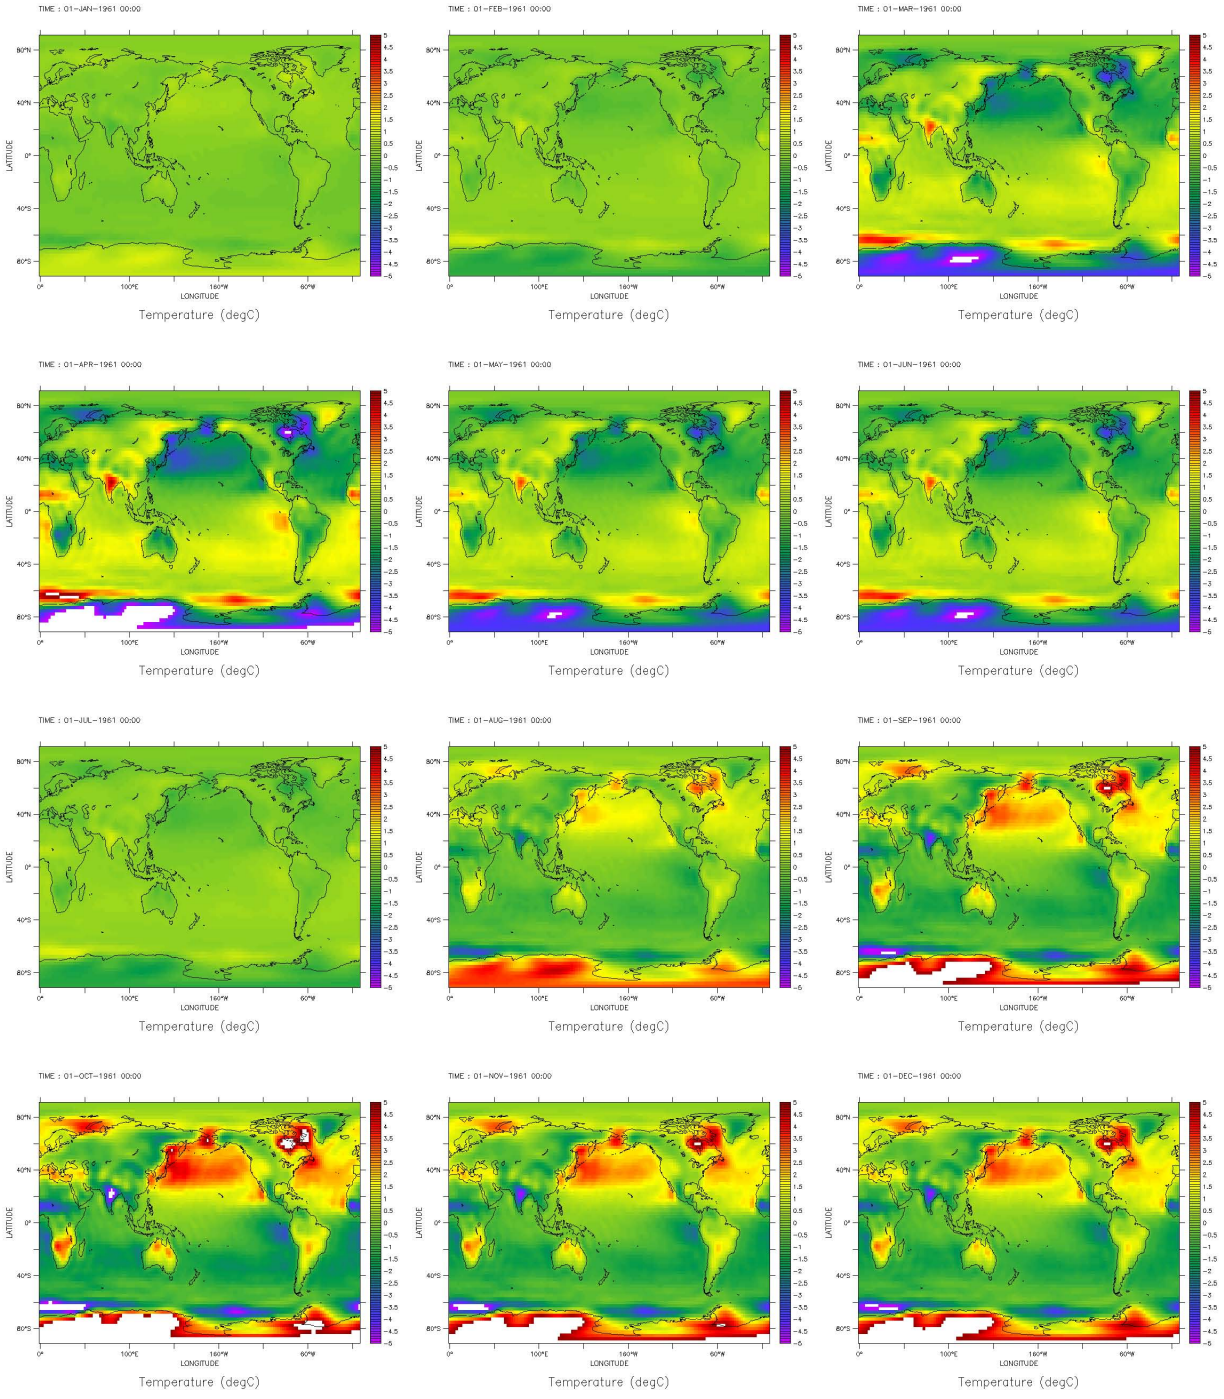

Figure 21: Spatial maps of IC 2 from the second Experiment and for the year 1961 (top left to bottom right).

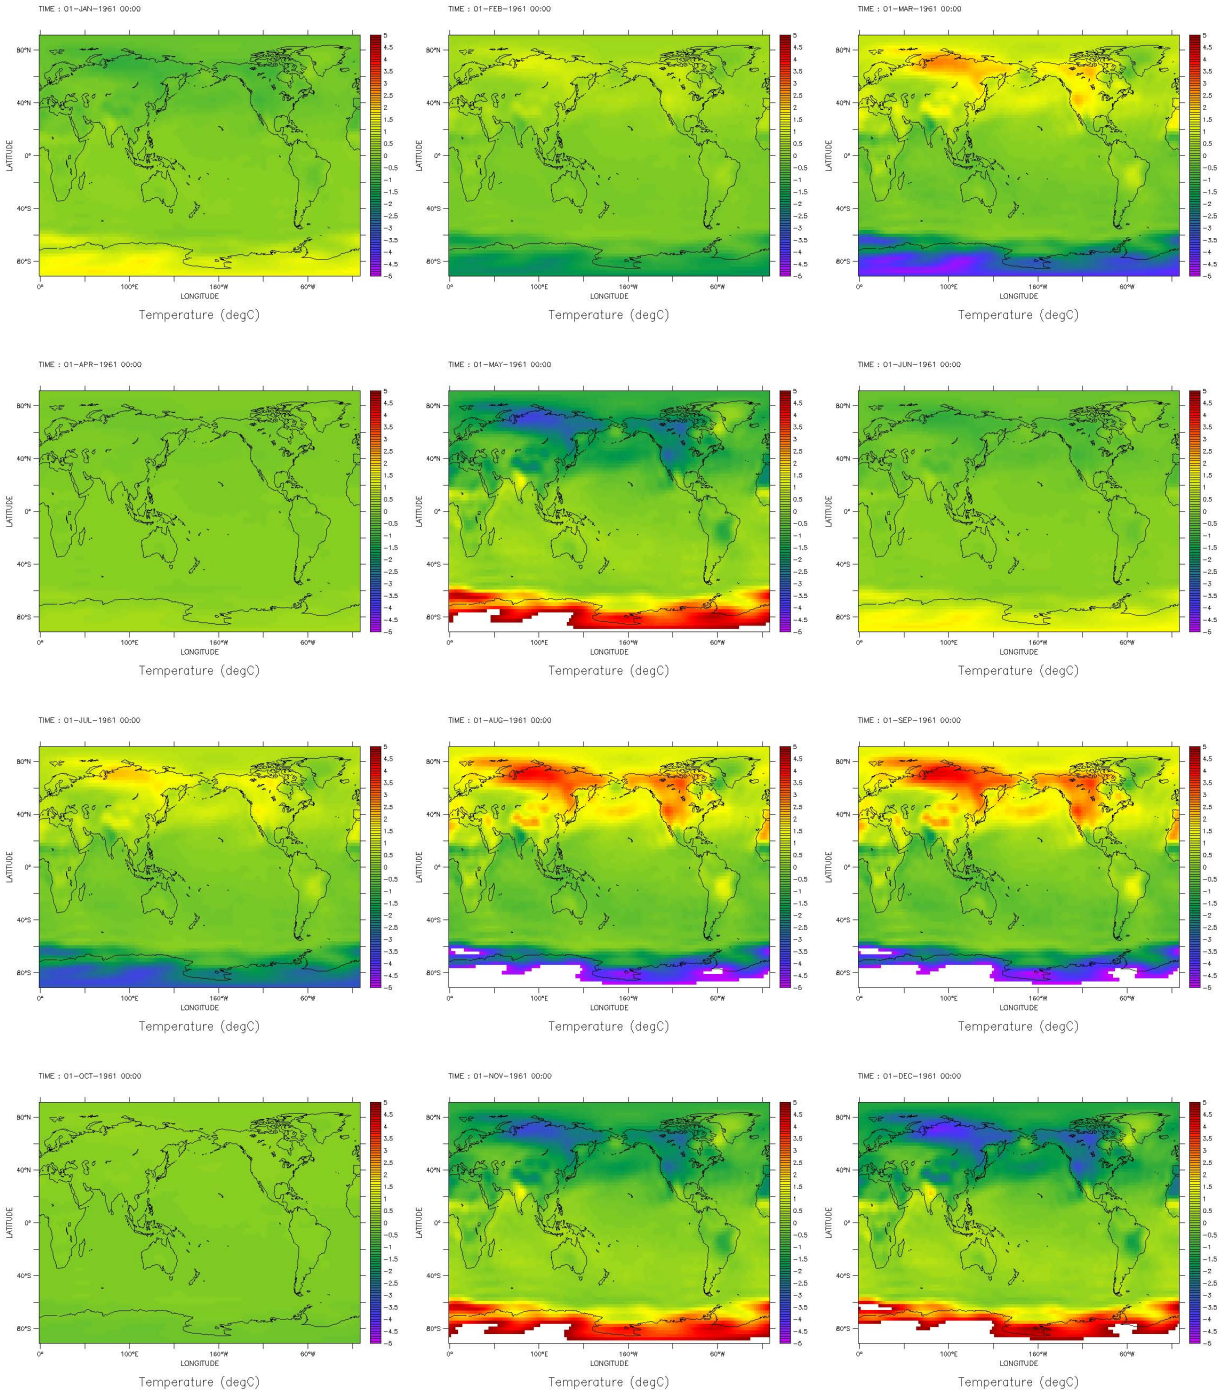

Figure 22: Spatial maps of IC 4 from the second Experiment and for the year 1961 (top left to bottom right).

4.4.3 North Atlantic Oscillation

The North Atlantic Oscillation (NAO) as by Burroughs and Williams (2003, p55), is a climate process which is characterised by a low pressure over Iceland and a high pressure over the Azores. When both the pressure systems are strong, the NAO is defined as being in its positive phase, and when they are weaker, it is defined as being in its negative phase. Its importance is attributed to the Westerly wind it defines, which during the positive phase creates an air flow that warms parts of Europe and Russia while cooling parts of Greenland. During its negative phase, it shows the opposite affects. The NAO climate index is the normalized pressure difference between a stations in Iceland and the Azores . We use the extended version provided by Jones et al. (1997)[5].

Hurrell et al. (2003) shows that negative NAO phases occurred in the years 1936, 1940, 1969 and 1977, and positive phases occurred in 1989 and 1990. As the data we have used overlaps with these time periods, we examine the spatial maps for the negative phase in 1969 generated from the different signal set sizes.

[5]NAO index: http://www.cru.uea.ac.uk/cru/data/nao/

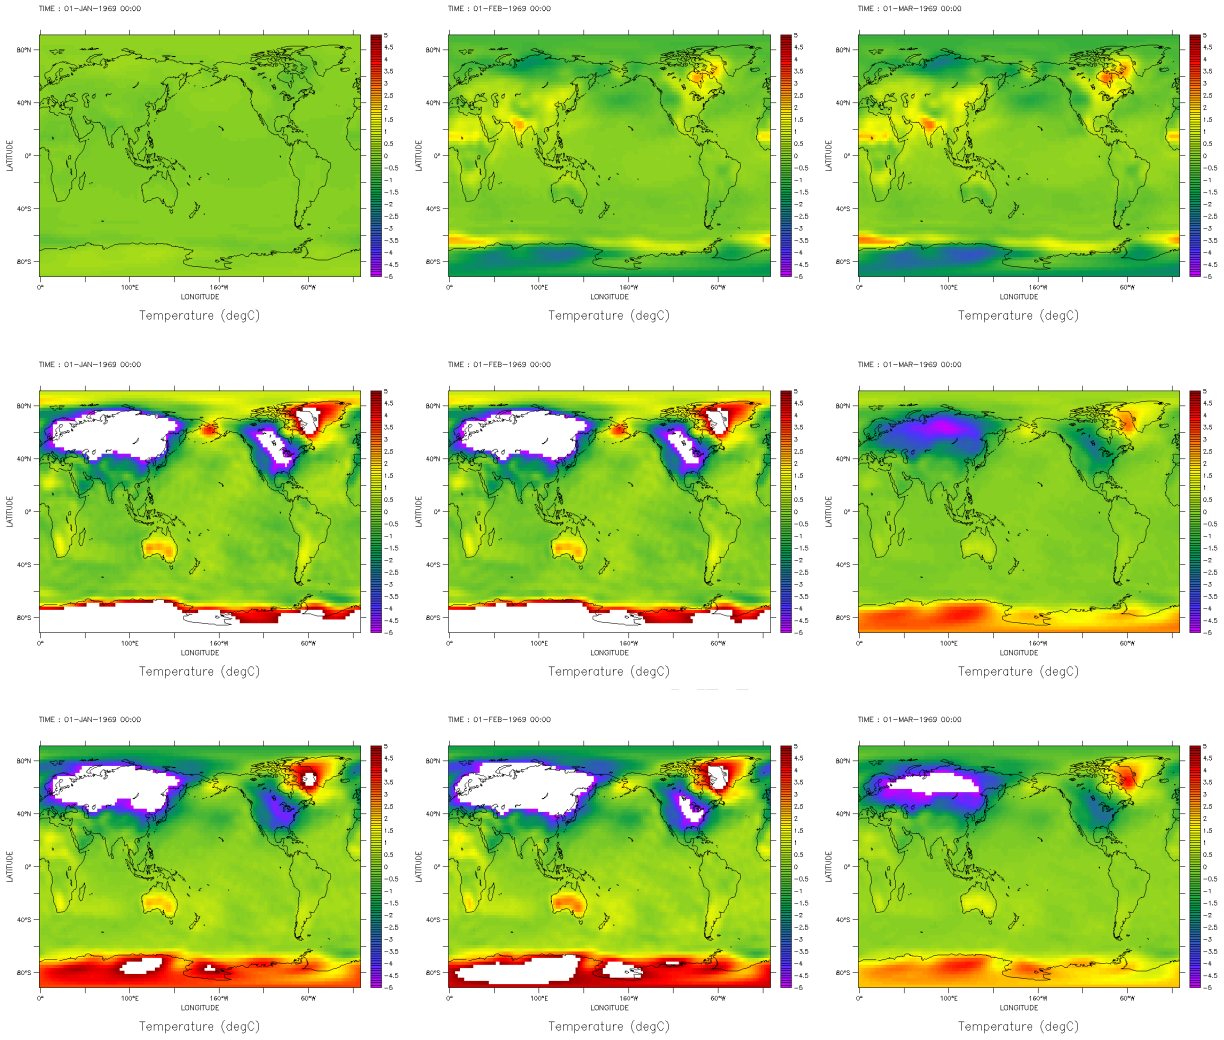

Figure 23: A negative phase of NAO showing the spatial maps for JFM of 1969. [*top*] Signal 4 from the set of 4 [*middle*] Signal 6 from the set of 6 [*bottom*] Signal 7 from the set of 10 as per table 7. White regions are beyond the temperature range of -5 to 5 °C.

Following the correlation of the signals to the popular climate indices, the spatial maps of the best matching signals were then examined. The spatial maps in figure 23 show the last 3 months of the negative phase of NAO in 1969. The best matching signals from the sets of 6 and 10 independent components show strong anomalies over Baffin Bay, Northern Asia, and North America. The main differences between the signals from the different sets are that the spatial patterns vary in strength.

The first row of images in figure 23 shows the spatial maps for signal 4 from the set of 4 independent components. The spatial maps do not show the strong patterns visible in the sets of 6 and 10 signals, and additionally the absolute correlation value is lower than them. The reason for the weak patterns and the low absolute correlation value may be due to the process being split over other components.

4.4.4 El Niño Southern Oscillation

Burroughs and Williams (2003, p141-158) describe the El Niño Southern Oscillation (ENSO) as the interaction of the Southern Oscillation and SSTs across the tropical Pacific. The El Niño and La Niña events are deviations from the climate norm of ENSO and indicate periods of SST warming and cooling in the tropical Pacific respectively. The importance of ENSO is that El Niño and La Niña events are able to have global consequences through its interaction with the Southern Oscillation. The index we have chosen to use is the Niño 3.4 index by Trenberth and Stepaniak (2001)[6].

A known period when ENSO was active was during the years of 1997 and 1998 when an El Niño event occurred. The period to examine the El Niño event was chosen as October, November, and December (OND) of 1997 as this is the time the spatial maps were most likely to represent ENSO.

Although the El Niño event occurred during the years of 1997 and 1998, the spatial maps for OND in figure 24 from the set of 4 and 10 signals do not show many strong patterns. However signal 3 from the set of 6 shows patterns over Northern Asia and North America. This is interesting as the correlation value for signal 8 from the set of 10 was the highest out of all the sets. The spatial maps have not presented any evidence of a temperature change over the central Pacific Ocean, and therefore the linking of the components to ENSO is difficult other than by using the comparison to the index. Westra et al. (2010) showed similar results when comparing 4 independent components against the Niño 3.4 index. They have showed that multiple independent components can represent different aspects of ENSO, and in their work the signals represent the variability of ENSO in the tropics and extra-tropics.

[6]Niño 3.4 Index: http://www.cgd.ucar.edu/cas/catalog/climind/TNI_N34/index.html#Sec5

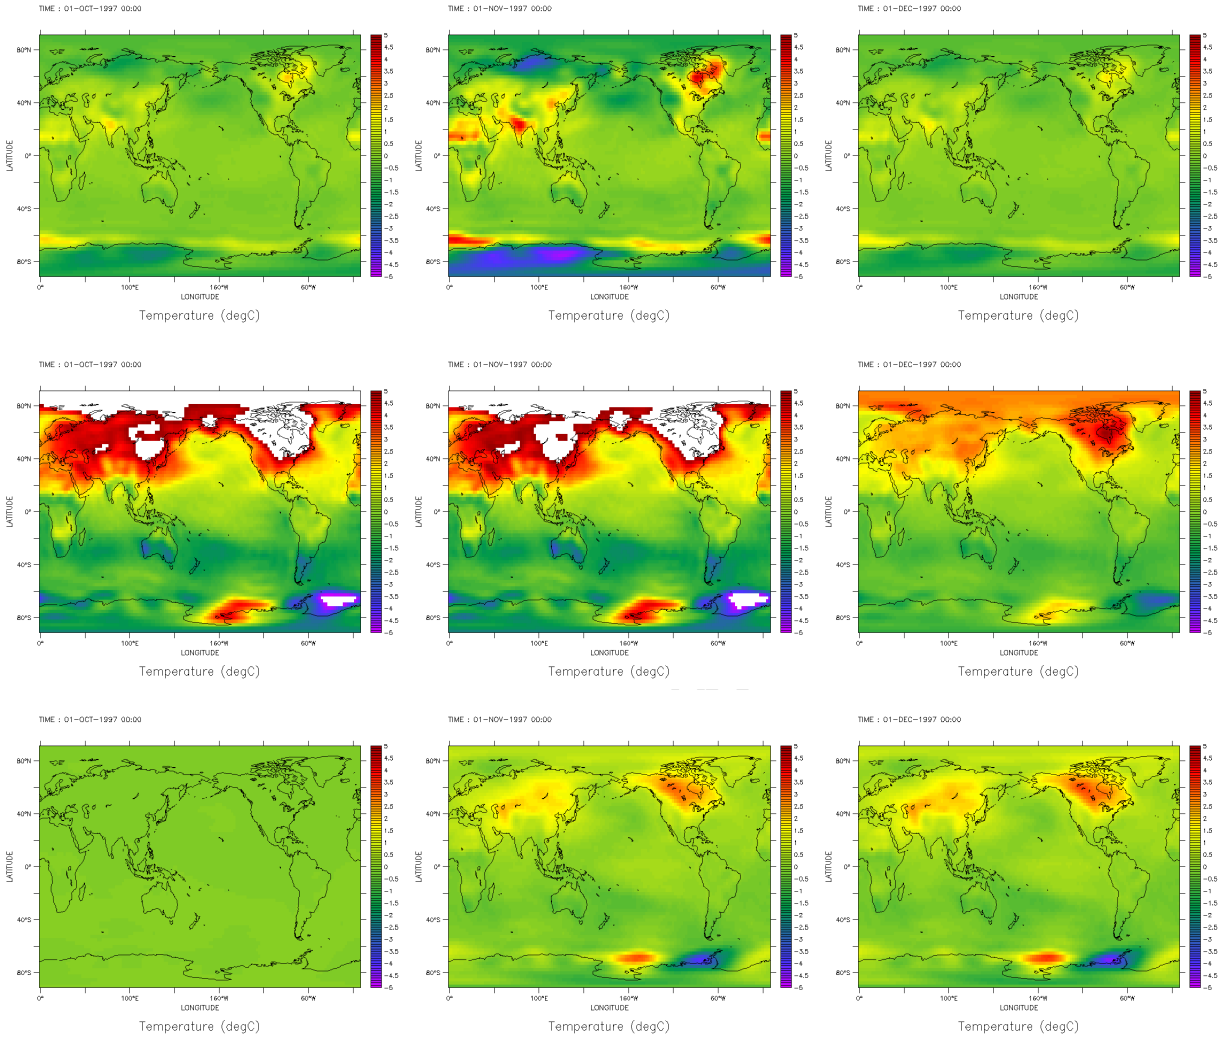

Figure 24: The spatial maps of OND for 1997 during part of an El Niño event. [*top*] Signal 4 from the set of 4 [*middle*] Signal 3 from the set of 6 [*bottom*] Signal 8 from the set of 10 as per table 7. White regions are beyond the temperature range of -5 to 5 °C.

4.4.5 Antarctic Oscillation

The Antarctic Oscillation (AAO) as defined by Gong and Wang (1999), is the difference of zonal mean sea level pressure between 40°S and 65°S. The derivation of the index was based on the application of PCA to their specified NCEP data. Contrary to this Marshall (2003) has argued that the NCEP data the index was derived from, had errors at high latitudes. To solve this they developed a benchmark from which they calculate a new index to represent the AAO. It is this index[7] that we use to compare our signals to.

Carvalho et al. (2005) derived phases of the AAO, and showed that negative phases are associated with El Niño and positive phases with La Niña events. One such event was the La Niña event from 1999 to 2000, which occurred during the months of December, January, and February (DJF). It is during this period that we examine our results to determine if our spatial maps show any pronounced patterns.

The spatial maps from the experiments during the La Niña period of 1999-2000 are shown in figure 25. Signal 3 from the set of 4 shows a spatial pattern over the Antarctic during this period. However the spatial maps from the other signals do not show any spatial patterns. The reason for this may be due to the splitting of the AAO over multiple signals and so the correlations of the signals are low and their spatial maps poorly present the spatial patterns associated with the process.

[7] AAO index: http://www.nerc-bas.ac.uk/icd/gjma/sam.html

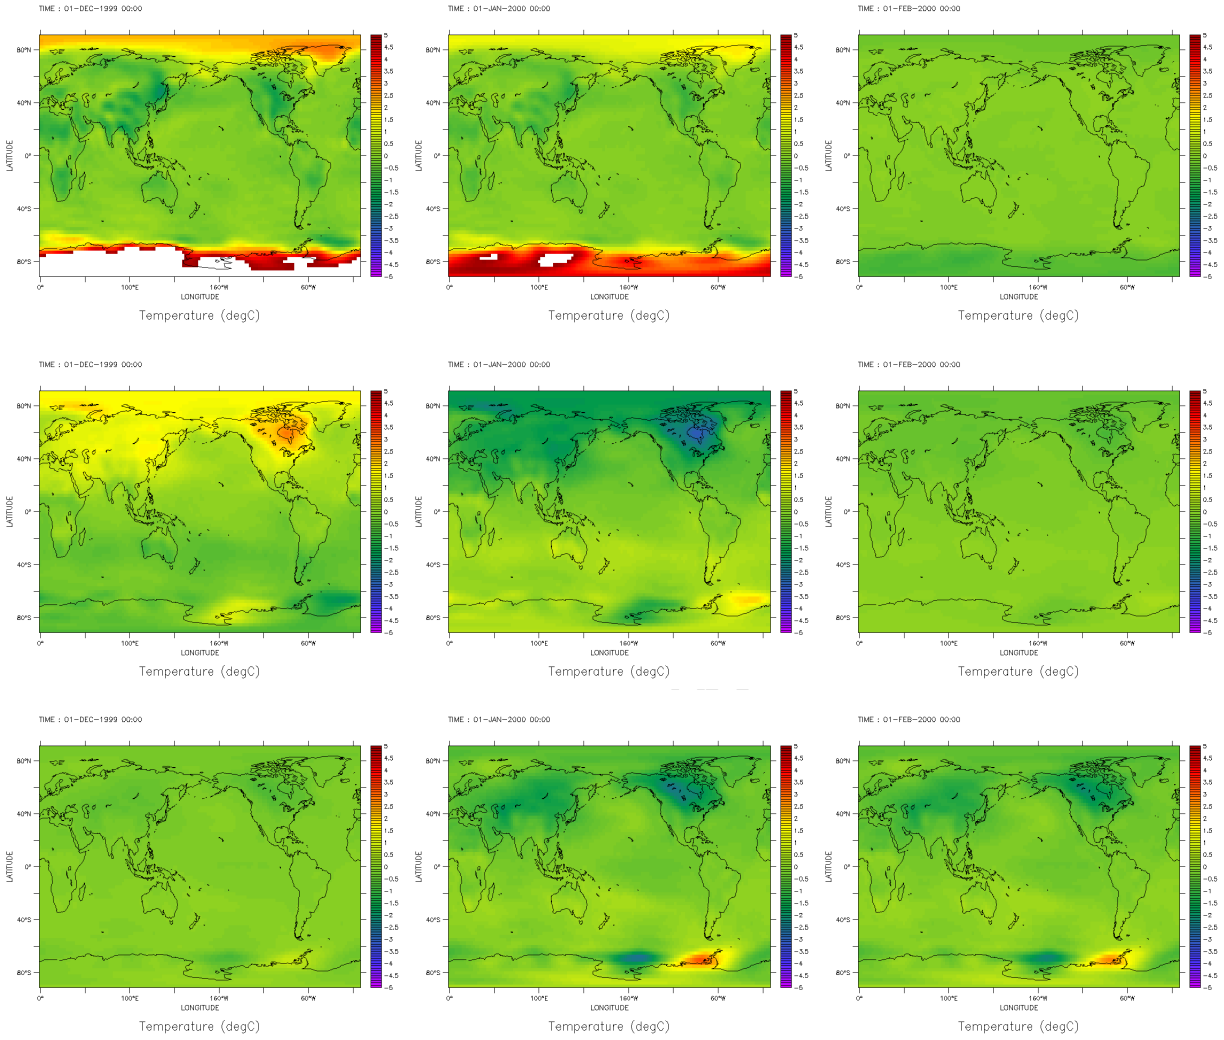

Figure 25: The spatial maps of DJF over 1999 to 2000 during part of a La Niña event. The spatial maps are from [*top*] Signal 3 of the set of 4. [*middle*] Signal 3 of the set of 6. [*bottom*] Signal 8 of the set of 10.

4.5 Conclusion

We have created a series of spatial maps which show the spatial manifestation of the separated signals over time. The spatial maps have been correlated against previously defined climate indices, and also compared to the climate processes already defined in the literature.

The first ranked signal from all the experiments represented the Seasonal Cycle, which subsequently accounted for a largest proportion of the variance (about 80%). A largely ocean based seasonal cycle was also identified as the second signal in two of the three experiments. In addition to this signals 4 and 5 of the set of 6 showed the effect of snow cover on the seasonal cycle, as the signals were separated from the land based seasonal cycle. This suggests that ICA is able to identify multiple different climate processes.

We found that the spatial maps best represented the North Atlantic Oscillation (NAO), as two of the best matched signals presented spatial patterns during a negative phase of the NAO. Their absolute correlations were also the highest for the NAO index. The spatial maps did show their limited use when two of the three best matched signals failed to show any meaningful spatial patterns during an El Niño event, and again when two of the three best matched signals failed to show spatial patterns of the AAO during a La Nina event.

One reason that ENSO was not well represented by the spatial maps may be due to the use of non-normalised data. In effect, the non-normalised data has the largest variability in the high latitudes, while the tropics have lower variability due to the dominance of the Seasonal Cycle in the higher latitudes. Therefore the non-normalised data may have been focused away from where ENSO would have been manifested within the spatial maps. Another reason that ENSO is not well represented, may be that the spatial maps may need a high correlation to the index in order to contain spatial patterns that are representative of it. The reason for the low correlation value may be due to the splitting of climate processes amongst several signals. The effect of this may be that the variance and the patterns associated with the process are also spread out over the signals. So some of the signals may better match the index, while other signals may contain the spatial patterns which are more representative of the climate process.

The spatial maps are also limited by the breadth of the literature. If no climate event can be found for the climate index in the literature, then there is no benefit to creating the spatial maps. Lastly as the focus of the literature is on specific events when creating spatial maps, then any events lasting more than a few months are difficult to analyse using only a small number of spatial maps. We have also assumed that the climate processes will always be manifested within our temperature data, which may not have been correct to assume for the AAO.

Despite these limitations the spatial maps may rather be considered as another option to examine signals with *low* correlations to previously defined indices. This is because the spatial maps can provide the spatial patterns of known events which may aid in the linking of a signal to a climate process when the temporal correlation to the index failed to do. This may be beneficial when the climate process is split over multiple signals, with each signal only weakly matching the index. Lastly, as ICA was able to detect the shift in the use of NCEP data in 1979 it may also be valuable for detecting other data inhomogeneities.

5 Conclusions

The aim of this work was to determine the value of Independent Component Analysis (ICA) by examining the ability of the spatial manifestations of its results over time, to identify climate processes (Eg: seasonal cycle) associated with the variability of the atmosphere. This work builds onto the methodology established by Wallace and Gutzler (1981). Whereby the understanding of a process is moved from pure theory, to a process which is supported by evidence in data. This increases the support for both the understanding of the process and for the method which was used to identify it.

The existing methods to identify climate processes have been successful in examining them within data and have created a number of innovations to improve the robustness of the results and the interpretations of them. We introduced ICA as an different method which created non-Gaussian and statistically independent patterns (signals) with which climate process could be identified.

5.1 Sensitivity of the Signals

One fundamental assumption of the climate processes is that the larger the percentage of variance explained (PVE), the more important the process is. This is because the larger the PVE by a climate process the more variability of the atmosphere it can be used to predict. Our contribution was to look at the sensitivity of the signals in terms of their PVE when we changed the number of signals separated, and when we used random values as the initial estimate of the signals.

We have shown that the PVE by the signals changes with a change in the number of signals separated from the data, with the first ranked (by PVE) signal gradually dropping in its PVE over a number of different sets. Westra et al. (2010) also showed that the results of ICA are sensitive to the number of signals estimated to be in their data.

In mapping the similarity of the signals over a number of runs when using a random initial estimate of them, we found that different signals are more structurally stable than others. The prime example of this was the first signal in our experiments which was always stable. The second ranked signal was found to be stable, but this was only the case in the set of 10 signals. However the remaining signals varied in their stability, which further indicated that there was no direct relationship between PVE by the signals and their stability. This suggests that within our data, after the first ranked signal there is no direct relationship between the PVE values and the stability of the remaining signals. Although the first solution to the stability results of the signals may appear to be in using a fixed initial estimate, the other signals generated by the initial random estimates may be just as equally valid. This raises the new question of how to analise different results which are all equally valid.

There were some short comings with the analysis. Firstly, neither additional data sets nor time periods were considered to perform ICA on and therefore the results presented are limited by the data and the time period chosen. Secondly, as during the preprocessing steps some of the variance of the data was discarded, additional information on some of the signals may have been lost which would have aided in our analysis.

5.2 Effect on Spatial Maps

Westra et al. (2010) showed that the El Niño Southern Oscillation was represented by not one but rather many signals in their work. This makes the attribution of the signals to the climate process difficult, especially when the signals only weakly correlate to previously defined climate indices that are associated with the climate process. Our contribution was to construct a series of spatial maps to aid in the linking of signals to climate processes.

We found that the first ranked signal from all the experiments represented the seasonal cycle over the land, while the second ranked signal represented the seasonal cycle focused mainly over the ocean. The ocean based seasonal cycle showed a two month shift behind the land based seasonal cycle which is most likely due to the thermal inertia of the oceans (Meehl et al., 2007). Out of the remaining climate processes, the spatial maps best represented patterns for the North Atlantic Oscillation. However they failed to represent both the El Niño Southern Oscillation and the Antarctic Oscillation, with only one of the three best index matching signals from different sets showing spatial patterns at the time of an event. However, the latter is not surprising as it may not have a large temperature manifestation.

The effect on the spatial maps when using a different number of signals resulted in both in different correlation values to indices and in different spatial patterns presented within the spatial maps. This indicates that the spatial maps were sensitive to changes in the number of signals separated. In terms of the sensitivity of the spatial maps to changing the initial estimate of the signals, this was never explicitly examined. However by examining the PVE changes observed from the mapping stability section, it is proposed that the changes in the signals would result in different correlations to the indices and that the spatial patterns within the spatial maps would also change.

Although not a direct effect on the spatial maps, the data change of NCEP to using full satellite data in 1979 was reflected in two signals from different sets. This may show an additional value in using ICA as it was able to detect data inhomogeneities. The actual effect of the data change was not further investigated as it was assumed to be negligible in our work.

A limitation of the spatial maps was that we used non-normalised data which emphasised the high latitude regions and could have masked the tropical regions where ENSO could have been seen in the spatial maps. A second limitation was that the identification of the event in the spatial maps relied on there being documented events, and that the signals would always capture them in their spatial maps. This may not have always been true to assume. However, we have show that ICA can be valuable in identifying both climate processes such as the ocean based seasonal cycle and in identifying data inhomogeneities.

5.3 The Limits of Independent Component Analysis

This section provides a summary of the comparison between ICA and other methods that could be used to identify climate processes within data. It reviews the advantages of ICA compared to other methods, and also examines the implementation problems that occurred.

The aim of using ICA was to separate signals from data, that were both maximally non-Gaussian and independent of each other. By constraining the signals to be independent, the signals would only represent a single climate process. In doing so, ICA would be able to overcome the mixing of climate processes within a single component, that other identification methods experienced. For example, both PCA and RPCA have been shown in the literature, to mix climate processes within their extracted components. The benefit to using non-Gaussian signals is shown by Aires et al. (2000), who identified ENSO within their SST data using ICA. Therefore this shows that climate processes with non-Gaussian distributions can exist in climate data. Similarly, correlation maps and graph communities only examine lower order correlations, and therefore they cannot maximise independence between non-Gaussian signals. The consequence of this is that both maximally independent and non-Gaussian signals could not be found using former methods.

While ICA has a number of potential advantages when identifying climate processes, it still suffers from a number of implementation problems. One of these problems was that the signals were found to be highly sensitive to a change in the number of mixtures used (section 3.2). Even the addition of one mixture could substantially change the distribution of the variance amongst the independent components. This problem is not seen in PCA, where although deciding the number of components to keep is still dependent upon the context, the resulting components are not effected by the number of them that are retained.

Another problem experienced by ICA, is the dependency of the results upon the initial estimate of the unmixing matrix (section 3.3). The concern is that we do not know if the spatial patterns identified in a signal will still be present if we use another initial estimate. Without quantifying the dependency, this can make it difficult to determine the effect on the association methods. However, this is not a problem for correlation maps, where the results are robust. Furthermore, the dependency of the results on the FastICA algorithm were not fully determined, and another implementation of ICA may have produced different results.

Lastly, during the association of signals to climate processes, ICA was shown to produce multiple signals which could be associated to an individual climate process. This was advantageous in that they were not associated to multiple climate processes, and therefore ICA was not mixing climate processes. However, with ICA splitting some of the climate processes over multiple signals, an analysis problem similar to the mixing problem was seen. Interestingly, this may suggest the existence of an ideal mixture size, whereby one signal represents one climate process, neither mixing nor splitting climate processes. Additional research will be needed to mitigate the problems of ICA, and refine its role in identifying climate processes within data.

5.4 Future Work

Perhaps one further application of ICA could be to use it to solve part of the multi-model problem. The current problem with using multiple models is outlined by Tebaldi and Knutti (2007). They state that models may be simulating the climate correctly but through the wrong methods. The first method is that the models may be trained and tested on the same data and therefore they could appear artificially good at simulating the present day climate. The second method they suggest is that the models may share components rather than have unique components. In sharing the components some biases within the components may not be canceled out and they can therefore be passed on to be reflected within the projections. The last method is through an unrealistic equilibrium being established between climate processes from different models. The climate processes may be incorrectly simulated by the models, but due to the cancellation of their errors amongst other models, the results are wrongly made correct.

We propose that ICA could be used to better understand the climate processes represented across models. As we have shown in our work the seasonal cycle can be separated from data using ICA, thus it may also be possible to extract the seasonal cycle from multiple models. As a start to solving the problem, we could work backwards and using the mixing matrix, determine which models best represented the seasonal cycle. The models could then be ranked according to their ability to represent the seasonal cycle. The importance of the seasonal cycle is shown by Knutti et al. (2006), where they use it to determine the performance of the models according to their sensitivity. As an additional aspect the use of the ocean seasonal cycle separated by ICA may also provide insights into the behavior of climate processes and feedbacks over the ocean.

6 References

F. Aires, A. Chédin, and J.P. Nadal. Independent Component Analysis of multivariate time series: Application to the tropical SST variability. *Journal of Geophysical Research*, 105:17, 2000.

A.D. Back and T.P. Trappenberg. Input variable selection using Independent Component Analysis. In *Proceedings of International Joint Conference on Neural Networks*, volume 2, pages 989–992. Citeseer, 1999.

A.G. Barnston and R.E. Livezey. Classification, seasonality and persistence of low-frequency atmospheric circulation patterns. *Monthly Weather Review*, 115(6):1083–1126, 1987.

J. Basak, A. Sudarshan, D. Trivedi, and MS Santhanam. Weather data mining using Independent Component Analysis. *The Journal of Machine Learning Research*, 5:239–253, 2004.

A.J. Bell and T.J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.

W.J. Burroughs and GE Williams. *Weather cycles: real or imaginary?* Cambridge University Press Cambridge, 2003.

L.M.V. Carvalho, C. Jones, and T. Ambrizzi. Opposite phases of the Antarctic Oscillation and relationships with intraseasonal to interannual activity in the Tropics during the austral summer. *Journal of Climate*, 18:702–718, 2005.

B. Christiansen. Atmospheric circulation regimes: Can cluster analysis provide the number? *Journal of Climate*, 20:2229–2250, 2007.

R.H. Compagnucci and M.B. Richman. Can Principal Component Analysis provide atmospheric circulation or teleconnection patterns? *International Journal of Climatology*, 28(6):703–726, 2008.

D. Dommenget and M. Latif. A cautionary note on the interpretation of EOFs. *Journal of Climate*, 15(2): 216–225, 2002.

I.K. Fodor and C. Kamath. Using Independent Component Analysis to separate signals in climate data. In *Proceedings of the SPIE*, volume 5102, pages 25–36, 2003.

D. Gong and S. Wang. Definition of Antarctic oscillation index. *Geophysical research letters*, 26(4):459–462, 1999.

L. Goulet and J.P. Duvel. A new approach to detect and characterize intermittent atmospheric oscillations: Application to the intraseasonal oscillation. *Journal of the Atmospheric Sciences*, 57:2397–2416, 2000.

J.W. Hurrell, Y. Kushnir, G. Ottersen, and M. Visbeck. An overview of the North Atlantic oscillation. *Geophysical Monograph-American Geophysical Union*, 134:1–36, 2003.

A. Hyvärinen. Survey on Independent Component Analysis. *Neural Computing Surveys*, 2(4):94–128, 1999.

A. Hyvärinen and E. Oja. A fast fixed-point algorithm for Independent Component Analysis. *Neural computation*, 9(7):1483–1492, 1997.

A. Hyvärinen and E. Oja. Independent Component Analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000.

I.T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, New York, 1986.

I.T. Jolliffe. A cautionary note on artificial examples of EOFs. *Journal of Climate*, 16(7):1084–1086, 2003.

PD Jones, T. Jonsson, and D. Wheeler. Extension to the North Atlantic Oscillation using early instrumental pressure observations from Gibraltar and South-West Iceland. *International Journal of Climatology*, 17 (13):1433–1450, 1997.

E.C. Kalnay, M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, et al. The NCEP/NCAR 40-year reanalysis project. *Bulletin of the American Meteorological Society*, 77(3):437–471, 1996.

R. Knutti, G.A. Meehl, M.R. Allen, and D.A. Stainforth. Constraining climate sensitivity from the seasonal cycle in surface temperature. *Journal of Climate*, 19:4224–4233, 2006.

I. Koch and K. Naito. Dimension selection for feature selection and dimension reduction with Principal and Independent Component Analysis. *Neural computation*, 19(2):513–545, 2007.

W.J. Krzanowski. Between-group comparison of Principal Components - some sampling results. *Journal of Statistical Computation and Simulation*, 15(2):141–154, 1982.

A. Lotsch, MA Friedl, and J. Pinzon. Spatio-temporal deconvolution of NDVI image sequences using Independent Component Analysis. *IEEE Transactions on Geoscience and Remote sensing*, 41(12 Part 2):2938–2942, 2003.

G.J. Marshall. Trends in the Southern Annular Mode from observations and reanalyses. *Journal of Climate*, 16:4134–4143, 2003.

G.A. Meehl, T.F. Stocker, W.D. Collins, P. Friedlingstein, A.T. Gaye, J.M. Gregory, A. Kitoh, R. Knutti, J.M. Murphy, A. Noda, S.C.B. Raper, I.G. Watterson, A.J. Weaver, and Z.C. Zhao. 2007: Global Climate Projections. In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. 2007.

J.P. Nadal, E. Korutcheva, and F. Aires. Blind source separation in the presence of weak source. *Arxiv preprint cond-mat/0005258*, 2000.

J.E. Overland and RW Preisendorfer. A significance test for Principal Components applied to a cyclone climatology. *Monthly Weather Review*, 110(1):1–4, 1982.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2010. URL `http://www.R-project.org`. ISBN 3-900051-07-0.

D.A. Randall, R.A. Wood, S. Bony, R. Colman, T. Fichefet, J. Fyfe, V. Kattsov, A. Pitman, J. Shukla, J. Srinivasan, R.J. Stouffer, A. Sumi, and K.E. Taylor. Climate models and their evaluation in: Climate change 2007: The physical science basis. contribution of working group i to the fourth assessment report of the intergovernmental panel on climate change. 2007.

M. Scholz, S. Gatzek, A. Sterling, O. Fiehn, and J. Selbig. Metabolite fingerprinting: detecting biological features by Independent Component Analysis. *Bioinformatics*, 2004. Figerprinting,.

M. Steinbach, P.N. Tan, V. Kumar, S. Klooster, and C. Potter. Discovery of climate indices using clustering. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, page 455. ACM, 2003.

K. Steinhaeuser, N.V. Chawla, and A.R. Ganguly. An exploration of climate data using complex networks. In *Proceedings of the Third International Workshop on Knowledge Discovery from Sensor Data*, pages 23–31. ACM, 2009.

J.V. Stone. *Independent Component Analysis: a tutorial introduction*. Mit Press, 2004.

C. Tebaldi and R. Knutti. The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365 (1857):2053–2075, 2007.

K.E. Trenberth and D.P. Stepaniak. Indices of El Niño evolution. *Journal of Climate*, 14:1697–1701, 2001.

H. von Storch and A. Navarra. *Analysis of climate variability: applications of statistical techniques*. Berlin, 2nd edition, 1999. ISBN 3-540-66315-0.

H. von Storch and F.W. Zwiers. *Statistical analysis in climate research*. Cambridge University Press, 1999.

J.M. Wallace and D.S. Gutzler. Teleconnections in the geopotential height field during the Northern Hemisphere winter. *Monthly Weather Review*, 109(4):784–812, 1981.

F. Westad and M. Kermit. Cross validation and uncertainty estimates in Independent Component Analysis. *Analytica Chimica Acta*, 490(1-2):341–354, 2003.

S. Westra, C. Brown, U. Lall, I. Koch, and A. Sharma. Interpreting variability in global SST data using Independent Component Analysis and Principal Component Analysis. *International Journal of Climatology*, 30(3):333–346, 2010.