University of Reading

School of Mathematics, Meteorology & Physics

Univariate Aspects of Covariance Modelling within Operational Atmospheric Data Assimilation

Nicholas Bird

August 2010

This dissertation is submitted to the Department of Mathematics in partial ful lment of the requirements for the degree of Master of Science

Declaration

I con rm that this is my own work, and the use of all material from other sources has been properly and fully acknowledged.

Nicholas Bird

Abstract

Operational weather centres around the world make use of data assimilation in order to combine observations with a forecast model to produce an estimate of the state of the atmosphere. The accuracy of this estimate is dependent upon a prior (`background') estimate of the atmospheric state.

The background state will always be subject to errors and these errors are specified using the Background Error Covariance Matrix, denoted by B. This matrix is almost impossible to model explicitly, so an approximation may be formed using a Control Variable Transform (CVT). The CVT uses a combination of parameter and spatial transforms in order to construct this approximation to the B matrix.

A number of investigations are performed on the spatial transform used in two operational centres. A two-dimensional version of the spatial transform is constructed using a combination of eigendecompositions and Discrete Fourier Transforms, using data supplied by the Met Office from an operational weather forecast model. This data consists of a set of forecast-differences on a single latitude ring.

It is shown that the correlations stored in the approximation are non-separable for

Contents

List of Notation

$\mathbf{B}$ Background error covariance matrix.

$\mathbf{B}^{imp}$ Implied background error covariance matrix.

$\mathbf{x}$ State vector.

$\mathbf{x}_t$ `True' state vector.

$\mathbf{x}_b$

List of Acronyms

NWP Numerical Weather Prediction

ECMWF European Centre for Medium-Range Weather Forecasts

CVT Control Variable Transform

3D-Var Three-dimensional Variational Assimilation

PDF Probability Density Function

FFT Fast Fourier Transform

DFT Discrete Fourier Transform

BLUE Best Linear Unbiased Estimate

NMC National Meteorological Centre

FFTW Fastest Fourier Transform in the West

LAPACK Linear Algebra PACKage

List of Figures

Chapter 1

Introduction

The aim of Numerical Weather Prediction (NWP) is to provide a best estimate of the true state of the atmosphere that is as accurate as possible. This estimate of the true state, referred to as the model state, can then be used in a forecast model in order to predict the future state of the atmosphere. An accurate prediction of the future state allows accurate weather forecasts to be made by NWP operational centres, such as the Met O ce and the European Centre for Medium Range Weather Forecasts (ECMWF).

1.1 Motivation

The use of a forecast model to produce a future model state can be seen as an initial value problem (Kalnay (2003)) and as a result requires a set of initial conditions in order to nd a solution. Observations of the atmospheric state can be used in order to provide these initial conditions, but often the number of observations available is not enough to provide a true representation of the whole atmosphere (Ide et al. (1997)). This results in an ill-posed problem (Bouttier and Courtier (1999)).

In order to ensure that the problem becomes well-posed the missing information in the initial conditions is provided by a prior estimate of the model state, which is commonly referred to as the background or rst-guess state (Bouttier and Courtier (1999)). This background state is generally produced through a previous run of the

forecast model (Ingleby (2001), Bannister (2008a)).

1.2 Aims of the Project

The aims of this project are to construct simplified models of the Met Office and ECMWF spatial control variable transforms, which will then be used to investigate the resulting $\mathbf{B}^{imp}$ matrix, aiming to solve the following questions:

1. What are the statistical and physical properties resulting from the Met Office and ECMWF spatial transforms?
2. What benefit does the ECMWF gain from the increased amount of data required to perform the spatial transform?
3. How important is dynamical and physical reasoning in constructing the vertical transforms used in the spatial transforms?

1.3 Main Results

The main results shown by this project are

1. The resulting covariances in the $\mathbf{B}^{imp}$ matrix resulting from both the Met Office and ECMWF spatial transforms are spatially homogeneous in the horizontal direction.
2. The correlation functions resulting from the spatial transforms are non-separable. Both the horizontal and vertical length scales increase with respect to vertical level, seen by observing the correlation between levels for a selection of vertical levels.
3. The higher magnitude in the leading vertical modes for the Met Office spatial transform is spread throughout the lower vertical levels in the ECMWF spatial transform.
4. The implementation of a mass-weighted outer product adds physical and dynamical reasoning into the spatial transform. No impact on the resulting implied B

matrix was seen, but an improvement in the power spectrum was noted for higher vertical modes.

1.4 Outline of Thesis

In Chapter 2 a description of Data Assimilation is given, focusing on three-dimensional variational assimilation (3D-Var). The Background Error Covariance Matrix is introduced, describing the role it has within data assimilation and why it is important. The structure of B is also described and some mathematical properties are presented. The Chapter also details different ways in which the B matrix is modelled. Particular focus given to the CVT method, since it is this method that is most relevant to the work done in the project.

Chapter 3 details the Met Office and ECMWF spatial transforms. In doing so the vertical and horizontal transforms are described in a way that is not specific to either operational centre. The main differences between the two implementations can then be seen, since the two centres construct the spatial transforms in different ways. The Chapter also introduces the `delta' test and a method of incorporating dynamical and physical reasoning into the spatial transform.

In Chapter 4 a number of experiments are conducted in order to answer the main questions set out in Section 1.2. Conclusions are made in Chapter 5 and the possibilities of further work are also explored.

Chapter 2

Data Assimilation and the Control Variable Transform

This chapter introduces the background material for this project. The concept of the Data Assimilation will be introduced, firstly in the general sense before focusing on the problem in terms of three-dimensional variational assimilation (3D-Var).

The chapter will then discuss the Background Error Covariance Matrix, which is the main focus of the project. Analysis of the structure of the Background Error Covariance Matrix will be performed, along with how it is modelled. The importance of it to the Data Assimilation process will also be discussed. The section ends with a discussion on how the Background Error Covariance Matrix can be decomposed into a form which will be of use in the project.

The final section of this chapter will outline the Control Variable Transform (CVT) method. The way in which it is introduced into the 3D-Var method will be discussed along with details on how it is performed in operational centres. The section will also

produce an analysis, denoted by $x_{a;k} \in R^n$ which represents the best estimate of the system.

The observations of $x_{t;k}$ valid at the current time are denoted by the vector $y_k \in R^p$, where p is the total number of observations made. Typically $p << n$, since fully observing the atmosphere is an almost impossible task (Bouttier and Courtier (1999), Ide et al. (1997)). The observations can be related to the model state through use of the observation operator H. H maps $R^n \to R^p$ and could be either linear or non-linear. This can be written as

$$y_k = H(x_{t;k}) + \epsilon_{y;k}; \qquad \epsilon_{y;k} \sim N(0; R) \tag{2.2}$$

$\epsilon_{y;k}$ represents the error in equation (2.2). This error could arise from the instruments used to make the observations, from the observation operator or through observation noise (Ide et al. (1997)). The observation error covariance matrix, $R \in R^{p \times p}$, stores the covariances of the observation errors. It is normally assumed that errors between observations are uncorrelated (Purser et al. (2003a)) and this assumption reduces R to a diagonal matrix, where the diagonal entries store the variances of the observations.

Since the length of the observation vector is generally smaller than the length of the model state vector, the initial conditions required for the forecast model need additional information in order for the problem to be well-posed (Bouttier and Courtier (1999)). This missing information provided by a prior estimate of the model state and is commonly referred to as the `background' state and is denoted $x_b \in R^n$. This background state is assimilated with the observation information to produce the required initial conditions.

The background state will also contain errors, and these errors are generally assumed to be taken from a Gaussian distribution with zero mean and covariance B. These errors, $\epsilon_{b;k}$, can be de ned by

$$x_{t;k} = x_{b;k} + \epsilon_{b;k}: \qquad \epsilon_{b;k} \sim N(0; B) \tag{2.3}$$

B is the background error covariance matrix and stores the covariances between the elements of x_b. A detailed description of B is given in Section 2.2. The assumption of

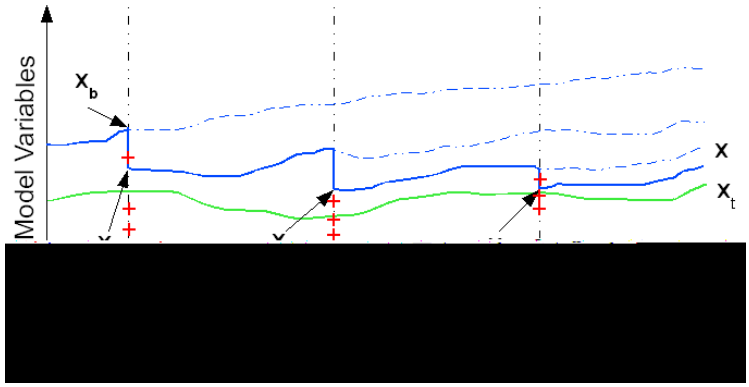

Figure 2.1: An illustration of how the process of finding the analyses of a system can be used over a number of time windows. For further explanation see text.

When the minimum of the cost function is found at the end of an assimilation window, the value x_a that is calculated becomes the starting point for the next assimilation window. By doing this, a more accurate estimate of the model state should occur as the process is repeated. The background state is taken to be the result of the forecast model at the end of the assimilation window. The observations for a given assimilation window are taken to be valid at the the same point as the background state for that window. Observations outside of the given assimilation window are not considered in this step.

2.1.3 Finding the Solution to the Cost Function

The objective of 3D-Var is to find the point at which the cost function in equation (2.4) is minimised. This point is the analysis and occurs when $x = x_a$. This point can be found by using gradient descent such as the Conjugate Gradient Method or the Method of Steepest Descent. Descriptions of the methods mentioned can be found in Iserles (2009). These methods, among others, look to find the minimum of the cost function by the use of iterative methods.

The minimum of the cost function is obtained when the gradient of the cost function, $\nabla J(x) = 0$ (Bouttier and Courtier (1999)). Using the assumption that H can be

linearised about x_b such that $H(x) \quad H(x_b) = H(x \quad x_b)$ (Courtier et al. (1994)), where H denotes the linear observation operator, the gradient is given by

$$r J(x) = B^{-1}(x \quad x_b) \quad H^T R^{-1}(y \quad H[x_b]): \tag{2.5}$$

For the case where the observation operator H is a linear operator, then the solution

biased estimate is applicable here since N is suitably large. It should also be noted here that this is the mathematical expectation for a finite sample of errors, which will result in an approximation to the $\mathbf{B}$ matrix.

In most operational centres equation (2.7) is simplified by assuming that that background errors are unbiased (i.e. $\langle \mathbf{b} \rangle = 0$)(Bannister (2008a), Lorenc (2003a)), which reduces the equation for the $\mathbf{B}$ matrix to

$$\mathbf{B} = \langle \mathbf{b}\,\mathbf{b}^T \rangle. \tag{2.9}$$

This demonstrates that for the unbiased case the $\mathbf{B}$ matrix is formed by taking the outer product (Anton and Rorres (2000)) on $\mathbb{R}^n$ of the background error vectors. By using this outer product definition the resulting $\mathbf{B}$ matrix will have important properties.

2.2.2 Properties of the B Matrix

The $\mathbf{B}$ matrix formed using equation (2.9) will be both square and symmetric with non-negative entries on the diagonal. This follows from the definition of the outer product and is also discussed in Section 6.1 of Anton and Rorres (2000). This is the case for all covariance matrices, not just those related to background error (Gaspari and Cohn (1999)). This property will be of relevance when discussing eigenvalues and eigenvectors in Section 2.2.7.

$\mathbf{B}$ is a positive semi-definite matrix since for any vector $\mathbf{v} \in \mathbb{R}^n$ (Hohn (1973)),

$$\mathbf{v}^T \mathbf{B} \mathbf{v} \geq 0; \tag{2.10}$$

$$\sum_{i=1}^{n} \sum_{j=1}^{n} \mathbf{v}_{(j)}^2 \mathbf{B}_{(i,j)} \geq 0. \tag{2.11}$$

2.2.3 Difficulties with Constructing the B Matrix

The `True' State

Perhaps the most important problem with constructing the $\mathbf{B}$ matrix using $\mathbf{b}$ is that the definition of $\mathbf{b}$ in equation (2.3) is reliant on the vector $\mathbf{x}_t$ (Bannister (2008a)). This

is problematic, since if the true state was known then the process of trying to estimate the truth would not be needed.

Number of Errors Required

By constructing B in the manner described in equation (2.9), the errors represented in the background of the system are modelled e ectively. Unfortunately, due to the large length of the state vector, actually creating this matrix is infeasible. In most operational centres $n \quad O(10^7 \quad 10^8)$ (Courtier (1997), Lorenc (2003a)). In order to ensure that B is of full rank, it is necessary to have at least n independent errors. In practice however the number of errors available B is likely to be much less than n. The e ects of using a B matrix that is not of full rank are highlighted in the paper by Hamill et al. (2000).

2.2.4 Methods of Constructing the B Matrix

The lack of the `true' state can be overcome through various methods, some of which are discussed by Fisher (2003) and Bannister (2008a). These can be divided into di erent classes of method, two of which are discussed here.

Using Surrogates of Background Error

One such method of constructing the B matrix involves using a surrogate of the background error in place of the true values. A popular method of doing this is the `NMC' method, named after the National Meteorological Centre (Fisher (2003), Parrish and Derber (1992)), although other methods do exist.

The NMC method uses di erences between runs of the forecast model to construct the surrogates of the background error. These forecasts will be of di erent length, with a typical di erence of 24 hours (Fisher (2003)), but the forecasts are chosen so that they nish at the same time. The di erence between the two forecasts form the background errors used to construct B (Bannister (2008a).

The main advantage of the NMC method is that the forecast data required to calculate the di erences is readily available in operational archives (Fisher (2003)). This

removes the computational expense required to perform the required number of forecast runs.

The NMC method does have disadvantages, which are discussed by Fisher (2003),

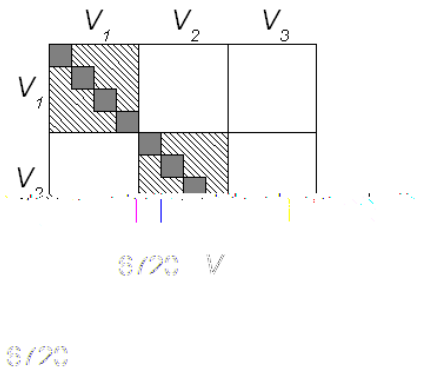

Figure 2.2: Example showing the structure of B in the case where x is prescribed from three variables on a two-dimensional grid. See text for more information.

Further de nition of these covariances can also be made, relating to the sub-matrix being considered. For element $B_{(i;j)}$, $i \neq j$, relating to the same variable, the element is an autocovariance (The hatched areas in Figure 2.2). For the case when $i = j$ the matrix entry is a variance as described above. The sub-matrices not on the block-diagonal represent cross-covariances or multivariate covariances, since they describe the covariance between two di erent variables (Bannister (2008a)).

2.2.6 The Role of B in Data Assimilation

The B matrix has an extremely important role in the data assimilation process. Some of the more important features of B are discussed in detail by Bannister (2008a) and are highlighted here to establish the need for a well-realised B matrix. If the B matrix does not represent the background error covariances well then it can have a detrimental impact on the resulting analysis.

1. Spreading of Observation Information

The rst such property is that the B matrix allows information from the observations to be spread from their position in grid point space to nearby points (Kalnay (2003), Bannister (2008a)). In order to demonstrate this property it is simpler to examine the

analysis increment resulting from the use of the BLUE as described in equation (2.6).

If a single observation is considered for this case, at grid point c for example, then it is clear that both y and R are reduced to scalar quantities denoted y_o and σ_o^2 respectively. The observation operator H will in this case represent a vector of length n. In this case H can be written as

$$H = \begin{cases} 1; & \text{if } i = c; \\ 0; & \text{otherwise.} \end{cases} \qquad i = 1;:::;n \tag{2.13}$$

By applying this H to the formula given in equation (2.6), the spreading of information can be demonstrated. The analysis increment at a position d, $d \neq c$, is given by

$$x_{a;d} \quad x_{b;d} = B_{(d;c)} \frac{y_o \quad x_{b;c}}{\sigma_o^2 + B_{(c;c)}}: \tag{2.14}$$

Here $x_{b;d}$ and $x_{b;c}$ are the values at positions d and c respectively of x_b. $B_{(d;c)}$ and $B_{(c;c)}$ are the entries of B at the respective positions. It is seen that information from the observed point has an effect on the other points, with the observation information weighted by the entry of B at point d.

The distance in which the observation information is spread is determined by the lengthscales of the correlations in B, which depends on the magnitude of the entries (Bannister (2008a)). This can be seen by performing single-observation experiments, such as those described in this subsection. It will also been seen in the application of the `delta' test, introduced in Section 3.4, which produces some of the results presented in Chapter 4.

The spreading of information is not limited to the univariate case, as the B matrix also allows for multivariate spreading (Bannister (2008a)). For example, this allows for observations of wind to affect other variables such as mass (Fisher (2003)). For this project only the univariate case in considered and this multivariate spreading will not be observed.

2.3 Control Variable Transforms

The method of constructing B that is of most interest in the scope of this project is that of using a Control Variable Transform (CVT). This is described in great detail by Bannister (2008b) and differences between operational centres are highlighted. Particular focus is given to both the Met Office and ECMWF (European Centre for Medium-Range Weather Forecasting) transform implementation, which ties in with the aims of this project.

2.3.1 Formulating the Control Variable Transform

Constructing B explicitly is almost impossible, due to the reasons outlined in Section 2.2. The aim of the CVT is to express the model variables in terms of new control variables. This is done with the aim of simplifying the background term of the cost function (see equation (2.4)). In order to explain how this works it is easier to express the cost function in it's incremental form (see Courtier et al. (1994)).

The incremental form of 3D-Var replaces the model state x in equation (2.4) by an incremental quantity $\mathbf{x} \in \mathbb{R}^n$. The model state increment is calculated from the model state x and a reference or `guess' state $\mathbf{x}_g \in \mathbb{R}^n$ such that

$$\mathbf{x} = \mathbf{x} - \mathbf{x}_g. \tag{2.17}$$

The background increment term is similarly defined as $\mathbf{x}_b \in \mathbb{R}^n$, where

$$\mathbf{x}_b = \mathbf{x}_b - \mathbf{x}_g. \tag{2.18}$$

Bannister (2008a) and Bouttier and Courtier (1999) note that the $\mathbf{x}_g$ is often taken to be equal to $\mathbf{x}_b$, resulting in a zero background increment. It is included here in order to show the relation between the two forms of 3D-Var.

In the incremental formulation the cost function is minimised with respect to $\mathbf{x}$ using the revised cost function (notation as defined by Ide et al. (1997)) as

$$J(\mathbf{x}; \mathbf{x}_g) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_b) + \frac{1}{2}\{\mathbf{y} - H(\mathbf{x}_g - \mathbf{x})\}^T \mathbf{R}^{-1}\{\mathbf{y} - H(\mathbf{x}_g - \mathbf{x})\}. \tag{2.19}$$

The other terms (B; R; y and *H*

This highlights the effectiveness of performing the CVT, namely that the background error covariance matrix is mapped to the identity matrix. This makes the problem more feasible to solve by removing the need to invert the B matrix. Once this more feasible problem has been solved the control variables can be transformed back into model variable terms in order to provide the analysis for the problem, which is given by

$$x_a = x_g + B^{1/2} \; : \tag{2.25}$$

The CVT as described works well if B can be split into its square root form, but as has been discussed in Section 2.2 the B matrix cannot be specified explicitly. Consequently it is not possible to express $B^{1/2}$ either. In order to alleviate this problem it is possible to construct an approximation of $B^{1/2}$ that models the B matrix. This approximation, denoted by $U \in \mathbb{R}^{n \times m}$ as described in Lorenc et al. (2000), is used to construct the implied B matrix B^{imp} such that

$$B^{imp} = UU^T. \tag{2.26}$$

The important task to perform when constructing the U matrix is to ensure that it captures the key properties that make up the B matrix itself. One of these properties is that B is symmetric and positive semi-definite. The B^{imp} matrix that is calculated from equation (2.26) will retain these features.

The effect of applying the CVT can be seen as a method of preconditioning the variational problem (Lorenc et al. (2000), Lewis et al. (2006)). The B matrix is likely to have a large range of eigenvalues, making it ill-conditioned. By performing the CVT, it is hoped that the problem becomes better conditioned.

2.3.2 Construction of CVTs using Parameter and Spatial Transforms

Control Variable Transforms in major operating centres are normally constructed by using a combination of parameter and spatial transformations, denoted by $K_p \in \mathbb{R}^{n \times m}$ and $B_s^{1/2} \in \mathbb{R}^{m \times m}$ respectively. The matrix U can then be written as

$$U = K_p B_s^{1/2}. \tag{2.27}$$

The notation used here is that used by Bannister (2008b) and is also detailed by Derber and Bouttier (1999).

The exact structure of these transforms varies from operational centre to operational centre. The pseudo-inverse of this equation will also be required (making the assumption that the pseudo-inverses exist), following from equation (2.23) and is

$$\mathbf{U}^{-1} = \mathbf{B}_s^{-1=2}\mathbf{K}_p^{-1}. \tag{2.28}$$

Discussion of the e ect of the parameter and spatial transforms will be done in terms of their pseudo-inverse, since this is applied in order to transform from the incremental variables into the control variables via equation (2.23).

The pseudo-inverse of the parameter transform $\mathbf{K}_p^{-1}$ works by transforming incremental variables from a given grid-space into new variables (denoted by $\overset{p}{1}$; $\overset{p}{2}$ etc.) that operate on the same grid-space. These new variables are stored in the vector $^p \in \mathbb{R}^m$ where

$$\mathbf{x} = \mathbf{K}_p^{-1}\ ^p. \tag{2.29}$$

The ^{p} variables are chosen such that they are assumed to be uncorrelated in order to allow them to be examined univariately (Parrish and Derber (1992), Lorenc et al. (2000)).

The pesudo-inverse of the spatial transform $\mathbf{B}_s^{-1=2}$ transforms the ^{p} variables into the control variables in using

$$= \mathbf{B}_s^{-1=2}\ ^p. \tag{2.30}$$

The spatial transform is used to remove any autocovariances between elds of a given parameter (Bannister (2008b)). This is illustrated in Figures 2.3 and 2.4.

Figure 2.3 shows the implied $\mathbf{B}$ matrix after the parameter transform has been applied in equation (2.29). The shaded segments represent areas where covariances between elements of the $\overset{p}{i}$ variables may exist. Each of the $\overset{p}{i}$ variables is itself a vector of length equal to the number of grid points. The white segments represent uncorrelated parameters and are therefore zero entries in the matrix. Although this

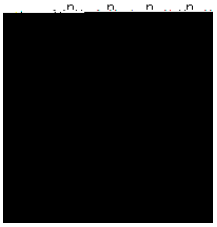

Figure 2.3: The implied B matrix after applying the parameter transform. See text for more information.

Figure 2.4: The implied B matrix after applying both the spatial and parameter transforms. See text for more information.

matrix is block diagonal, it is not yet in the form that is desired from performing the CVT.

Figure 2.4 shows the effect of applying both the spatial and parameter transforms. Each of the p_i parameters is projected onto spatial modes via the spatial transform. These modes are then normalised by dividing by the standard deviation of that mode. As the different spatial modes are uncorrelated this causes the spatially transformed parameters to be uncorrelated and have unit variance. The resulting matrix is the identity matrix and fulfils the aim of the CVT discussed in Section 2.3.1.

The way in which the different operating centres perform the parameter and spatial transforms differ. This may be due to the order in which certain steps are performed or even which parameters are chosen in the parameter transform stage. In Chapter 3 the spatial transforms for the Met Office and the ECMWF are discussed in more detail.

2.3.3 Assumptions made when performing CVTs

At the various stages of performing a Control Variable Transform there are a number of assumptions that are made. These assumptions have important consequences for the resulting $\mathbf{B}^{imp}$ matrix.

Possibly the most important assumption made in the CVT process are that the resulting horizontal correlation functions are homogeneous and isotropic (Buehner (2005)).

These assumptions are made as part of the spatial transform stage. Gaspari and Cohn (1999) and Hollingsworth and Lonnberg (1986) de ne these assumptions in detail, with further discussion given by Courtier et al. (1998). Since they are important features of the CVT and will have an impact on this project they are described here.

Homogeneity

A correlation function C_1(

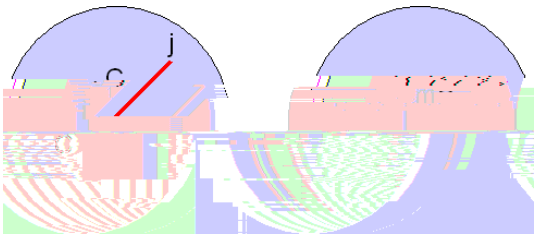

Figure 2.6: An example of an isotropic correlation function.

holds. An example of an isotropic correlation function is given in Figure 2.6.

This assumption of homogeneity and isotropy in the horizontal correlations is not an accurate assumption to apply in reality, since it can be shown that correlations are indeed inhomogeneous and anisotropic. Otte et al. (2001) discusses the limitations of applying isotropic error structures around observations. There has been a large amount of work into relaxing these assumptions through the use of various techniques. These techniques include the use of recursive filters (see Purser et al. (2003b) and Wu et al. (2002)), wavelet formulations (Fisher (2003)) and distorted grids (Segers et al. (2005), Desroziers (1997)).

2.4 Summary

Chapter 3

The Met Office and ECMWF Spatial Transform

In this chapter the methodology involved in performing the spatial transform stage of the CVT for two operational centres is presented. The methods described relate to the Met Office (Section 3.2) and ECMWF (Section 3.3) definitions of the spatial transforms, which are described in Bannister (2008b). In both of the methods presented in this chapter the data is prescribed on a two-dimensional grid of J horizontal grid-points over I vertical levels.

Some general concepts relating to the spatial transform are first introduced. These include the a recap of the aims of the CVT and the use of a calibration step in the spatial transforms. The form of the implied B matrix is described, linking this to the two methods that are to be used in this project.

The chapter then describes the two methods, beginning with the Met Office spatial transform before outlining the ECMWF spatial transform. The calibration steps required for each centre are detailed, before the algorithm for performing the spatial transform is presented. In presenting the methodology for both operational centres important differences between the two centres are highlighted.

Finally, an alternative outer product de nition of the vertical transform is presented that looks to incorporate physical and dynamical reasoning into the spatial transform that is otherwise not present.

3.1 General Concepts

3.1.1 Aims of the CVT

The CVT, which is introduced in Section 2.3, transforms model variables into a new set of control variables. These control variables form a representation that diagonalises the resulting implied Background Error Covariance Matrix, denoted by $\mathbf{B}^{imp} \in \mathbb{R}^{IJ \times IJ}$. The variable I denotes the number of vertical levels, with J denoting the number of horizontal grid-points. The aim is to construct a matrix $\mathbf{U}$

matrix is constructed using the vertical and horizontal transforms and involves the use of training data (Section 5 of Bannister (2008b)). This training data is used to produce some of the information that will be required in the course of calculating the vertical and horizontal transforms.

For the purposes of this project the univariate case is considered, with data supplied for a single control variable. This training data has been supplied by the Met O ce (see Chapter 4 for details of this data) and has already been transformed using the parameter transform described in Section 2.3. The variable being considered is the streamfunction ( ).

The Streamfunction

The streamfunction measures non-divergent ow upon a two-dimensional eld (Glickman (2000)). It can be related to the coordinate velocities u and v of the ow in the Cartesian $(x; y)$ plane by the following equations

$$u = \frac{@}{@y}; \tag{3.2}$$

$$v = \frac{@}{@x}: \tag{3.3}$$

Storing the Streamfunction Data

Let $\mathbf{D}$ denote the tensor storing the samples of streamfunction data. If the data is prescribed on a two-dimensional grid of J horizontal grid-points, I vertical levels and there are S samples of data in total, then $\mathbf{D} \in \mathbb{R}^{I \times J \times S}$. Individual element $\mathbf{D}_{(i;j;s)}$, corresponding to row i, column j of sample s, is given by

$$\mathbf{D}_{(i;j;s)}; \qquad \begin{aligned} i &= 1; : : : ; I; \\ j &= 1; : : : ; J; \\ s &= 1; : : : ; S: \end{aligned} \tag{3.4}$$

Using $\mathbf{D}$, the vertical covariance matrix $\mathbf{D}^{cov} \in \mathbb{R}^{I \times I}$ is calculated. Let $\mathbf{d}_{j;s} \in \mathbb{R}^{I}$ denote the vertical pro le of $\mathbf{D}$ corresponding to horizontal grid-point j of sample s.

Individual entries of $\mathbf{d}_{j;s}$

3.1.4 The Discrete Fourier Transform

In the spatial transforms described in this chapter, the horizontal transform U_H performs a one-dimensional transformation into spectral space. The transformation into spectral space is achieved through the use of a Discrete Fourier Transform (DFT) (see Iserles (2009)), which uses a Fast Fourier Transform (FFT) algorithm in order to extract the component frequencies relating to each horizontal wavenumber.

If a DFT is applied to a vector $v \in \mathbb{R}^J$ with elements $v_{(j)}$, $j = 1, \ldots, J$, the components of each individual horizontal wavenumber are extracted. This information is stored in a vector of length J with each element representing a horizontal wavenumber. The resulting vector $V \in \mathbb{C}^J$ has individual entries V

an inverse DFT to perform the transformation. The inverse DFT acts upon a vector of spectral coefficients $\mathbf{V} \in \mathbb{C}^J$. The result is a vector $\mathbf{v} \in \mathbb{R}^J$, with elements $v_{(j)}$ given by

$$v_{(j)} = \frac{1}{J} \sum_{n=1}^{J} V_{(n)} e^{\frac{2\pi i}{J}(j-1)(n-1)}, \qquad j = 1, \ldots, J \tag{3.12}$$

As with the standard DFT, this process can be written in matrix form and the inverse DFT matrix $\mathbf{W}^{-1} \in \mathbb{C}^{J \times J}$ such that $\mathbf{v} = \mathbf{W}^{-1}\mathbf{V}$ is defined as

$$\mathbf{W}^{-1} = \frac{1}{J} \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & \omega^{-1} & \omega^{-2} & \ldots & \omega^{-(J-1)} \\ 1 & \omega^{-2} & \omega^{-4} & \ldots & \omega^{-2(J-1)} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \omega^{-(J-1)} & \omega^{-2(J-1)} & \ldots & \omega^{(J-1)(J-1)} \end{pmatrix} \tag{3.13}$$

where $\omega = e^{-\frac{2\pi i}{J}}$ is again the nth root of unity.

The result of the DFT must be scaled in order to ensure that the inverse DFT is equal to the complex conjugate transpose of the DFT (Kincaid and Cheny (2002)). The DFT and inverse DFT described here do not model this requirement as they are not identical ($\mathbf{W}^{-1} \neq \mathbf{W}^*$, where * indicates the complex conjugate transpose). By scaling the DFT by a factor of $\frac{1}{J}$ the inverse DFT is equivalent to the complex conjugate transpose of the DFT.

The FFT used to perform the DFT and its inverse has the advantage that the number of operations required to perform it can be greatly reduced from that of calculating the DFT explicitly, provided that J can be expressed such that $J = 2^n, n \in \mathbb{Z}$ (Iserles (2009)). This makes it computationally efficient to use, since the calculation of the DFT and inverse DFT can be performed much more quickly than would be possible using the explicit calculation.

The DFT and inverse DFT are performed in the

3.2 The Met O ce Spatial Transform

3.2.1 Overview

In the Met O ce and ECMWF spatial transform the implied B matrix is formed by performing a combination of vertical and horizontal transformations. In the Met O ce spatial transform the spatial part of the matrix U is constructed using the matrix product $U_V U_H$, where $U_V \ 2 \ R^{IJ \times IJ}$ and $U_H \ 2 \ R^{IJ \times IJ}$. By calculating U using this product and the transpose U^T the form of the B^{imp} matrix for the Met O ce spatial transform is given by

$$\begin{aligned} B^{imp} &= UU^T; \\ &= U_V U_H U_H^T U_V^T: \end{aligned} \tag{3.14}$$

The vertical and horizontal transform matrices de ned in equation (3.14) will be re-de ned in Sections 3.2.2 and 3.2.3 to enable the calculation of the R matrix as described in Section 3.1.3. The steps presented in this section for the construction of the R matrix using the Met O ce spatial transform are discussed in Bannister (2008b) and Lorenc et al. (2000), with additional information provided by sta at the Met O ce (Wlasak (2010), Personal Communication).

3.2.2 Calibration

Before constructing the R matrix using the Met O ce Implementation, there are a number of calibration steps that are taken.

1. Vertical Transform Matrices

In order to construct the vertical transforms for the Met O ce spatial transform, it is necessary to perform the eigendecomposition of D^{cov} (see equation (3.7)). The eigende-composition was detailed in Section 2.2.7 and allows D^{cov} to be written in terms of its eigenvalues and eigenvectors. The eigenvalues are stored on the diagonal of the matrix $\ 2 \ R^{I \times I}$ and the eigenvectors stored in the matrix $E \ 2 \ R^{I \times I}$, where a column of E

stores a single eigenvector. The decomposition is written as

$$D^{cov} = E \quad E^T: \tag{3.15}$$

Since D^{cov} is symmetric and positive semi-de nite (a feature of all covariance matrices as discussed in Section 2.2) the eigenvalues in will be both non-negative and real-valued. If D^{cov}

2. Calculating the Power Spectrum

The next step in the calibration stage for the Met O ce spatial transform is to calculate the power spectrum of the sample data. The power spectrum will be used as part of the calculation of R and is a measure of the magnitude of the data attributed to each horizontal wavenumber.

The T_V matrix in the Met O ce spatial transform is used to transfer the data in D from full grid-point space into vertical mode space. The T_V matrix is post-multiplied by each of the *S* samples in the D tensor, to produce a tensor D^{T_V} *2* R

This operation removes the complex elements c ts and allows the power spectrum for an individual mode to be pl lex conjugacy property (see Section 3.1.4) of the horizontally tra ower spectra, only the rst $n + 1$ wavenumbers need to be plot the power spectrum for an individual mode. The then averaged over the S samples to produce the Individual entries are given by

$$\mathsf{D}^{P}_{(i;j)} = \frac{1}{S} \sum_{s=1}^{S} \mathsf{D}^{H_{mod}}_{(i;j;s)}; \qquad i \qquad j \tag{3.24}$$

Calculating D^{P} completes the calibration stage ial transform.

3.2.3 Calculating the R matrix

With the calibration steps completed, it is now pos matrix using the ordering of vertical and horizontal transforms 4).

1. The First Vertical Transform

The calculations that follow from here are assum on a matrix $\mathsf{M} \in \mathbb{R}^{I \times J}$, which could represent a matrix of sam ervation data at di erent grid-points. The matrix M ensures t the required dimensions. R will be constructed using the verti sform stages, along with the power spectrum calculated in Secti

This newly multiplied information is then transformed back from spectral space into vertical mode space, using the operator S^{-1}. This is a level-by-level inverse DFT (Section 3.1.4) on the result of the matrix product given by equation (3.27). The resulting vertical mode information is stored in the matrix U_H 2 $R^{I \times J}$, with entries

$$U_{H(i,:)} = W^{-1}[(D^P \quad U_H^T)_{(i,:)}]^T; \qquad i = 1;:::;I: \tag{3.29}$$

3. The Final Vertical Transform

The nal operation that is performed when calculating R is the second vertical transformation U_V, which was de ned in equation (3.16). This operation transforms back into full grid-point space. This is achieved by pre-multiplying the U_H matrix from equation (3.29) by U_V

for the construction of the R matrix using the ECMWF spatial transform are discussed in Bannister (2008b) and Derber and Bouttier (1999).

It should be noted that the ECMWF spatial transform described in this Section is

vector $\mathbf{L} \in \mathbb{R}^I$ which has entries defined by

$$\mathbf{L}_{(i)} = \frac{1}{J} \sum_{j=1}^{J} \mathbf{D}^{var}_{(i;j)}; \qquad i = 1; \ldots; I: \tag{3.34}$$

The entries of $\mathbf{L}$ are used to modify the data in $\mathbf{D}$, with this modified data stored in the three-dimensional tensor $\mathbf{D}$

The pseudo-correlation matrix $\mathbf{D}_j^{cor} \in \mathbb{C}^{I \times I}$ for wavenumber j, $j = 1,\ldots,J$, is formed by taking the covariance of the vertical profile of data corresponding to that wavenumber. This is performed over each of the S samples. Let $\mathbf{d}_{j,s} \in \mathbb{C}^I$ denote the vertical profile of data in $\mathbf{D}^H$ corresponding to sample s and wavenumber j. Then

$$\mathbf{D}_j^{cor} = \frac{1}{S} \sum_{s=1}^{S} (\mathbf{d}_{j,s} - \langle \mathbf{d}_{j,s} \rangle)(\mathbf{d}_{j,s} - \langle \mathbf{d}_{j,s} \rangle)^T; \qquad j = 1,\ldots,J. \tag{3.37}$$

In equation (3.37) the angled brackets denote mathematical expectation (see equation (2.8) in Section 2.2).

The expression given in equation (3.37) is simplified by making the assumption that the vertical profiles of data in $\mathbf{D}^H$ are unbiased (i.e. $\langle \mathbf{x}_{j,s} \rangle = 0$; Courtier et al. (1994), Bannister (2008a)). The pseudo-correlation matrix corresponding to wavenumber j is now given by

$$\mathbf{D}_j^{cor} = \frac{1}{S} \sum_{s=1}^{S} \mathbf{d}_{j,s} \mathbf{d}_{j,s}^T; \qquad j = 1,\ldots,J. \tag{3.38}$$

These matrices are referred to as pseudo-correlation matrices because they are not true correlation matrices in the standard sense. A correlation matrix (Gaspari and Cohn (1999)) is a covariance matrix which has been divided by its variance, whereas here the data has also been transformed into spectral space before the matrix is calculated. A true correlation matrix would have ones along the diagonal, which is not the case here.

By performing this step, one of the main differences between the two implementations covered has been highlighted. The ECMWF Implementation requires the storage of J pseudo-correlation matrices, each of size (I I). The Met Office Implementation only requires a single covariance matrix to be calculated which is of size (I I). If the values of I and J are sufficiently large this could result in a much greater amount of storage for the ECMWF Implementation over that required by the Met Office.

The final step in calculating the vertical transform matrices requires decomposing the pseudo-correlation matrices. As in the Met Office Implementation, an eigendecomposition is performed, extracting the eigenvalues and eigenvectors (see Section 2.2.7). This is performed on each of the $\mathbf{D}^{cor}$ matrices, resulting in $\in \mathbb{C}^{I \times I}$ and E $\in \mathbb{C}^{I \times I}$ matrices for each of the J wavenumbers. In the ECMWF spatial transform these are

denoted $_j$ and $\mathbf{E}_j$, where $_j; \mathbf{E}_j$ represent the matrices corresponding to wavenumber j.

The vertical transforms $\mathbf{U}_{Vj} \in \mathbb{C}^{I \times I}$ and $\mathbf{U}_{Vj}^T \in \mathbb{C}^{I \times I}$ can then be constructed, where $\mathbf{U}_{Vj}; \mathbf{U}_{Vj}^T$ denote the vertical transform matrices corresponding to wavenumber j. They are constructed using

$$\mathbf{U}_{Vj} = \mathbf{E}_j \, {}_j^{1/2}; \tag{3.39}$$

$$\mathbf{U}_{Vj}^T = {}_j^{1/2} \mathbf{E}_j^T. \tag{3.40}$$

In order for the $_j$ and $\mathbf{E}_j$ matrices to be of full rank, it is important that a sufficient number of samples are used to calculate the pseudo-correlation matrices. Since each matrix is calculated from a single of column of data per sample, there must be at least I samples to ensure full rank. If $S < I$ then zero-valued eigenvalues will appear in , which will have an impact on the subspace the vertical transform matrices can span (see Section 3.3 in Bannister (2008a)).

This requirement is not as important for the Met Office Implementation due to the way in which the covariance matrix is calculated from the sample data. As shown in equation (3.6) of Section 3.2.2, $\mathbf{D}^{cov}$ is calculated by summing across the J horizontal grid-points and also the S samples. Unless I is sufficiently large, then the equality $JS > I$ should hold and the matrix will be of full rank.

3.3.3 Calculating the R Matrix

With the calibration steps completed, it is now possible to calculate the R matrix using the ordering of vertical and horizontal transforms defined in equation (3.32). This involves the use of a matrix $\mathbf{M} \in \mathbb{R}^{I \times J}$ in the same manner as for the Met Office spatial transform. By using a matrix of these dimensions the resulting R matrix will be of the required dimensions.

1. Multiplication by Model Level Standard Deviations

The calculation of R proceeds by multiplying M by the standard deviation of each model level. If the standard deviations of the model levels are stored in a matrix $2\ \mathrm{R}^{I\times I}$, where is a diagonal matrix with

$$_{(i;i)} = \sqrt{\mathsf{L}_{(i)}}: \qquad i = 1;:::;I \tag{3.41}$$

Multiplying each element of M by the respective model level standard deviation is then calculated using the matrix product

$$\mathsf{M} \tag{3.42}$$

2. The First Horizontal Transform

The result of the matrix product in equation (3.42) is then transformed into spectral space using the operator *S*, which is a level-by-level DFT on the rows of M as described in Section 3.1.4. This transforms the data from full grid-point space into spectral space, allowing individual wavenumbers to be considered. The spectrally-transformed information is stored in the matrix $\mathsf{U}_H^T\ 2\ \mathrm{C}^{I\times J}$. Entries are given by

$$\mathsf{U}^T_{H(i;:)} = \mathsf{W}[(\ \mathsf{M})_{(i;:)}]^T; \qquad i = 1;:::;I: \tag{3.43}$$

3. The Vertical Transforms

The next stage involves performing the vertical transforms and is achieved by using the U_{Vj} and U^T_{Vj} matrices calculated in equations (3.39) and (3.40) in Section 3.3.2. This requires extracting the columns of U_H^T. Let $\mathsf{P}_j\ 2\ \mathrm{C}^I$ denote the jth column of U_H^T, $j = 1;:::;J$.

Each P_j is pre-multiplied rst by the U^T_{Vj} matrix corresponding to the corresponding wavenumber, then pre-multiplied by the corresponding U_{Vj} matrix. Let $\mathsf{P}^V_j\ 2\ \mathrm{C}^I$ denote these transformed columns, such that

$$\mathsf{P}^V_j = \mathsf{U}_{Vj}\mathsf{U}^T_{Vj}\mathsf{P}_j; \qquad j = 1;:::;J: \tag{3.44}$$

The P_j^V vector is the jth column in the matrix $U^P \in C^{I \times J}$ with entries

$$U^P_{(i;j)} = P^V_{j(i)}; \qquad \begin{array}{l} i = 1;\ldots;I; \\ j = 1;\ldots;J: \end{array} \tag{3.45}$$

In the vertical transform stage of the ECMWF spatial transform the data in the matrix U_H^T is transformed into vertical mode space and then back into spectral space.

4. The Second Horizontal Transform

The next step to be performed in the ECMWF spatial transform is the horizontal transform stage, denoted by the operator S^{-1}. S^{-1} is a level-by-level inverse DFT as described in Section 3.1.4 and transforms the rows of the matrix U^P back into full grid-point space. The resulting matrix is denoteed by $U_H \in R^{I \times J}$ and entries are given by

$$U_{H(i;:)} = W^{-1}(U^P_{(i;:)})^T; \qquad i = 1;\ldots;I: \tag{3.46}$$

5. Multiplication by Model Level Standard Deviations

The final step in the ECMWF spatial transform algorithm is a multiplication of the grid-points by the standard deviation on a level-by-level basis. This is a pre-multiplication of U^H given in equation (3.46) by . The result is the R matrix where

$$R = U_H: \tag{3.47}$$

R can also be written as

$$R = (S^{-1}[(U_V U_V^T)_j(S[\ M])]); \tag{3.48}$$

where $(U_V U_V^T)_j$ indicates the column by column vertical transform process described by equation (3.44). This completes the algorithm for calculating the R matrix using the ECMWF spatial transform.

3.4 The `delta' test

In order to investigate the Met O ce and ECMWF spatial transforms and observe some of their properties a `delta' test is performed. The `delta' test is a simpli cation of the single observation test which is used in many works to observe the analysis increment for a single observation (examples can be seen in Bouttier and Courtier (1999), Kalnay (2003) and Section 2.2.6 of this project). The simpli cation is that a single point in the eld is selected and the analysis increment calculated without any observation information.

This test involves choosing the matrix M as described in the descriptions of the implementations such that it is a basis matrix. This basis matrix 2 $\mathrm{R}^{I \times J}$ consists almost entirely of zeros. Only one entry in is non-zero and has the value one, in row m and column n. is therefore de ned by

$$_{(i;j)} = \begin{cases} 1 & \text{if } i = m; j = n; \\ 0 & \text{otherwise.} \end{cases} \tag{3.49}$$

Note that is not the standard Kronecker delta (Kincaid and Cheny (2002)) found

The `delta' test is then performed by using the following steps:

1. Set value of *n*.

2. For *m*

By implementing an alternative outer product (the mass-weighted outer product) the

(3.57) is valid since the inverse of $\mathbf{P}$ will always exist. No entry on the diagonal of $\mathbf{P}$ will be equal to zero since the p values will always be positive (there will always be a mass attributed to a vertical level).

With this information the vertical mode matrix $\mathbf{V} \in \mathbb{R}^{l \times l}$ can now be calculated. The columns of $\mathbf{V}$ store the vertical modes and $\mathbf{V}$ is calculated such that $\mathbf{V} = \mathbf{P}^{-2}\mathbf{Y}$. Using these vertical modes the vertical transforms $\mathbf{U}_V$, $\mathbf{U}_V^T$ and $\mathbf{T}_V$ are now defined as

$$\mathbf{U}_V = \mathbf{V}^{1/2}; \tag{3.58}$$

$$\mathbf{U}_V^T = \boldsymbol{\Lambda}^{1/2}\mathbf{V}^T; \tag{3.59}$$

$$\mathbf{T}_V = \boldsymbol{\Lambda}^{-1/2}\mathbf{V}^T\mathbf{P}^2. \tag{3.60}$$

These transforms can now be used in the construction of the $\mathbf{R}$ matrix as described in Sections 3.2 and 3.3.

3.6 Summary

In this chapter the spatial transforms used in the Met Office and ECMWF have been introduced. The two operational centres use a different combination of the vertical and horizontal transforms in order to construct the $\mathbf{B}^{imp}$ matrix. A number of calibration steps are performed before the construction of $\mathbf{B}^{imp}$ occurs, with the steps varying depending on the operational centre.

The vertical transform is constructed by performing an eigendecomposition, but is defined differently for each centre. In the Met Office the eigendecomposition is performed on a vertical covariance matrix constructed from sample data. In the ECMWF the vertical transform matrices are calculated on data which has had the model level variance removed and been transformed into spectral space. The vertical transform matrices are allowed to vary by horizontal wavenumber, increasing the storage required for the ECMWF spatial transform.

The ECMWF spatial transform also requires a greater number of samples in order to ensure that the vertical transform matrices are of full rank. This is not the case in the Met Office spatial transform, due to the way in which the vertical transform differs.

The horizontal transform is performed by using a DFT to transform into spectral space and an inverse DFT to transform out of spectral space.

The `delta' test that will be used in Chapter 4 was de ned, which is a simpli ed version of a single observation test. The `delta' test will be used order to test the assumption of homogeneity that is made during the spatial transform process. Finally, a description of the introduction of an alternative outer product in the construction of the vertical transforms was presented.

Chapter 4

Analysing the Implied B Matrix

In this Chapter an analysis of the implied B matrix that results from the Met O ce and ECMWF Spatial Transforms is presented. This analysis is performed in order to test some of the key assumptions that are made during the CVT process. The investigations that make up this analysis are designed to answer the main aims of the project as set out in Chapter 1. These aims will be met by trying to answer the following questions through the analysis of the implied B matrix:

1. What are the statistical and physical properties resulting from the Met O ce and ECMWF spatial transforms?
2. What bene t does the ECMWF gain from the increased amount of data required to perform the spatial transform?
3. How important is dynamical and physical reasoning in constructing the vertical transforms used in the spatial transforms?

The Chapter opens with a look at the sample data that was provided in order to conduct the investigations that were conducted. The results presented here were produced using this sample data for both the Met O ce and ECMWF spatial transforms. The results produced from conducting the `delta' test are then presented, which show the results of making the assumption of spatial homogeneity in the CVT.

Question 1 is answered by observing the correlations that are present between vertical levels, using the implied C matrix through a variant of the `delta' test. Comparisons between the two operational centres are drawn from these results.

Further comparison between the two centres is made by analysing the leading vertical modes in both cases. By doing this analysis some insight can be made into the extra storage requirements of the ECMWF, which relates to question 2.

Finally, some discussion is made regarding question 3, focusing on the implementation an alternative outer product into the spatial transform stage. This outer product makes changes to the vertical transform stage by including a scaling matrix before the eigendecomposition is performed. The impact this has on the spatial transform will be discussed.

4.1 The Sample Data

4.1.1 Description of Sample Data

In order to produce the output in this chapter a number of samples of data were provided by the Met O ce. The number of samples required in order to produce meaningful results varies depending on the implementation being considered (Section 3.3.2), however it is important that the data is at a consistent resolution.

The sample data used in this project was produced from forecast di erences taken from a standard Met O ce data assimilation experiment run for 28 days in June 2008, with the purpose of testing forecast accuracy. The experiment includes forecast model runs for 6 hour and 30 hour time periods, run at resolution N320, which nish at the same time. The resolution dictates the maximum number of wavenumbers that are used in the experiment and gives a measure of the number of grid points that the model state is prescribed upon.

To determine the number of grid-points that are used, consider the general case NX, where X is a numerical value. The grid upon which the latitude ring is de ned has

3/2X + 1 Meridional (North-South) Grid-Points running from the north to south pole (including the poles).

Since this project details a simpli ed two-dimensional grid the meridional grid points are not taken into account for the sample data used in this project.

The data used for this project is of a much lower resolution than the results of the experiment provide, so the full- elds are recon gured into N48 resolution. Once this is done, the di erence is taken between the two model forecasts in order to provide the data in the correct increments. This is an example of the NMC method for obtaining forecast statistics (see Section 2.2.4). The number of grid-points upon which the data is prescribed is dependent upon the resolution that is used.

Using an N48 resolution results in data that is prescribed on a two-dimensional grid which consists of 96 horizontal grid-points. Data is also prescribed upon 70 vertical model levels, encapsulating the troposphere to the mesosphere (Barry and Chorley (2010)). In Table 1 in the Appendix the model level heights and equivalent pressures are given for the 70 model levels.

The data samples provided for this project hold data for a single latitude ring, corresponding to the streamfunction control variable (see Section 3.1.2). 28 samples of data were produced for four latitude rings. These rings are at 45N, 45S, 40N and 40S. This gives a total of 112 samples that will be used in the experiments in this Chapter. The data for the four latitude rings is grouped together.

It is recognised that there are likely to be slight di erences between the latitude rings provided, but it is assumed that these di erences are negligible for the purposes of the project. The data supplied for the streamfunction has been chosen speci cally in order to avoid any signi cant di erences between degrees of latitude. It can be seen that at certain locations on the globe (i.e. around the Tropics) the data would have a signi cantly di erent shape (see Ingleby (2001)). For the latitude rings provided for it is assumed that the data will be suitable to use with each other in order to ensure that the ECMWF spatial transform is constructed using a large enough number of samples.

4.1.2 Examining the Grid-Point Variances

In order to interpret the results that will follow in later sections it is important to understand the shape of the data as it is given. To do this the grid-point variances are plotted. These grid point variances are stored in the matrix D^{var} calculated using equation (3.33) in Section 3.3.2.

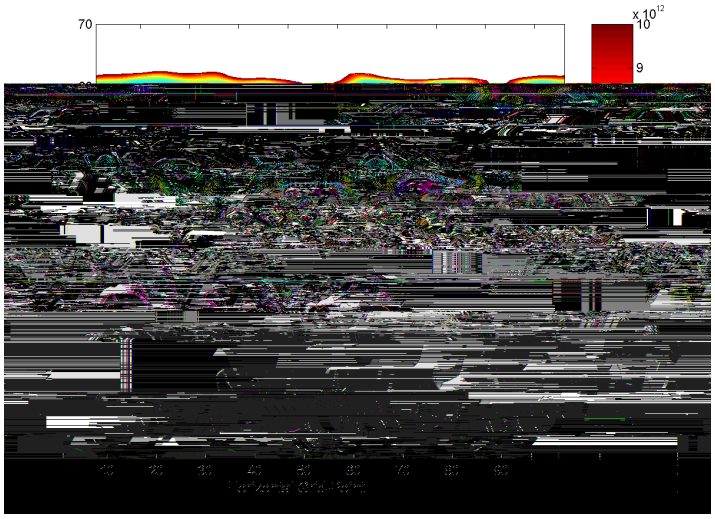

Figure 4.1: Plot of the variances of the sample data for the streamfunction. Data in the range (1 10^{11} to 10 10^{12}) is plotted.

Once the grid-point variances have been calculated the values can be plotted using a contour plot. Figure 4.1 shows the contour plot produced for the variable using equation (3.33) calculated from sample data. Only the information in the range from (1 10^{11} to 10 $10^{12} m^2/s$) has been plotted here in order to observe the shape around the lower vertical levels more clearly. Plotting the full range of the vertical levels makes it more difficult to see the information in the lower vertical levels, as the variance in the higher levels is of $O(10^{15}) m^2/s$.

As can be seen from the plot, there is an increase in the variances around vertical level 40, which then drops down again before beginning to rise toward the higher vertical levels. The increase in the variances around level 40 can be attributed to the presence

of a jet causing turbulence in the air around this level. The high variances in the top ve levels (representing the top of the stratosphere and the mesosphere) show a large increase in wind-speed. This is a known unphysical feature of the data, resulting from a lack of resolution in the upper levels of the forecast model as well as the upper boundary conditions of the model (Wlasak (2010), Personal Communication).

4.1.3 The Empirical Vertical Covariance Matrix

An important diagnostic tool for the analysis throughout this Chapter is the empirical vertical covariance matrix. This matrix is calculated in the calibration stages of the Met O ce and ECMWF Spatial Transforms in Sections 3.2 and 3.3 and denoted by D^{COV} in equation (3.7) of Section 3.1.2. Due to the low resolution being used for the sample data in this project it is possible to calculate this matrix and perform the eigendecomposition required to produced the information required for the vertical transform stages.

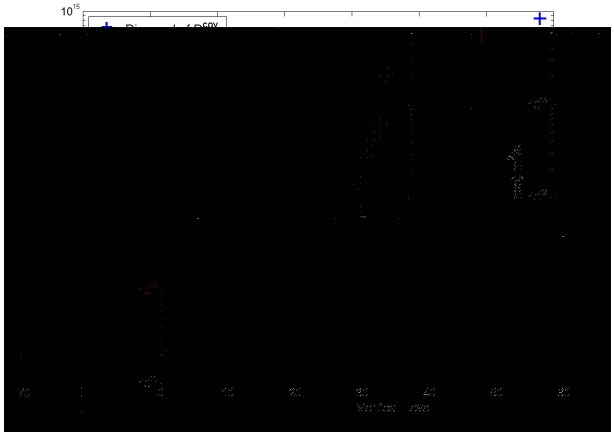

Figure 4.2: Plot in log-scale of the diagonal of the D^{COV} matrix.

Figure 4.2 is a plot of the diagonal of the empirical vertical covariance matrix. By observing the variance at each point on the diagonal, the information shown in the contour plot of grid-point variances (Figure 4.1) is observed more clearly. The variances is greatest in the top ve vertical levels, with the increase in variance at vertical level 40 also seen.

4.2 Results from the `delta' Test

This section will present the results obtained from performing the `delta' test, for the R matrix generated as result of both the Met O ce (see Section 3.2) and ECMWF (see Section 3.3) spatial transforms. Details of the `delta' test are given in Section 3.4.

In order to test the homogeneity assumption, this process is repeated over a number of columns. If the resulting columns are identical, then the data in R is spatially homogeneous. For the purposes of the experiments conducted in this Section, the `delta' test was performed over columns at horizontal grid-points 10, 50 and 80.

4.2.1 Results using the Met O ce Spatial Transform

Figure 4.3 shows the results that were obtained from producing the R matrix using the Met O ce Spatial transform. The column corresponding to horizontal grid-point 80 is plotted together with the diagonal of the D^{cov} matrix obtained from the data samples. This experiment shows that the transform procedure has worked correctly, as the two lines are virtually identical.

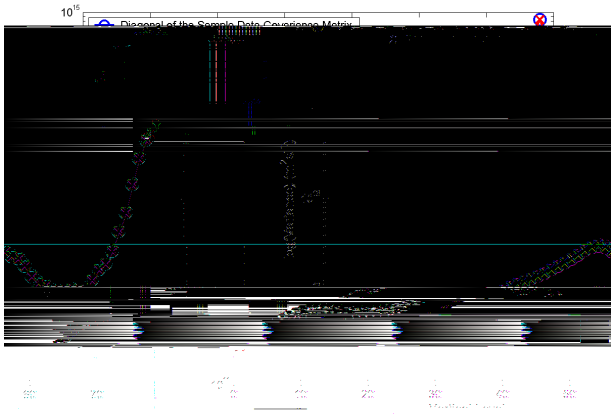

Figure 4.3: Plot in log-scale showing the results of the Met O ce Spatial Transform. Plotted are the diagonal of the D^{cov} matrix (blue line) and the column of R corresponding to horizontal grid-point 80 (red line).

From Figure 4.3, it can also be seen that the variance increases sharply from vertical level 55 onwards. This matches the shape of the sample data shown in Figure 4.1. The small increase at vertical level 40 is also featured in the results shown in Figure 4.3.

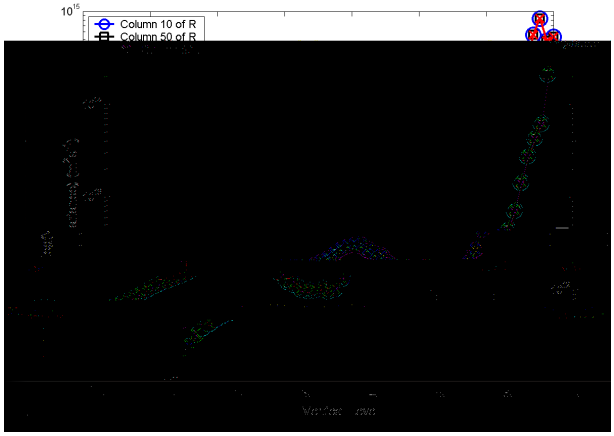

Figure 4.4: Plot in log-scale showing the three columns of R calculated using the `delta' test with the Met O ce Spatial Transform. The blue line corresponds to column 10. Column 50 is plotted in black and column 80 in red.

In order to examine the homogeneity assumption the three columns are plotted on the same graph, which is shown in Figure 4.4. Although slightly di cult to distinguish the three lines from each other, the resulting columns produced are almost identical. Investigation of the di erence between the lines shows that any di erence is of

the three columns together, shown in Figure 4.6.

Figure 4.5: Plot in log-scale showing the results of the ECMWF Spatial Transform. Plotted are the diagonal of the D^{cov} matrix (blue line) and the column of R corresponding to horizontal grid-point 80 (red line).

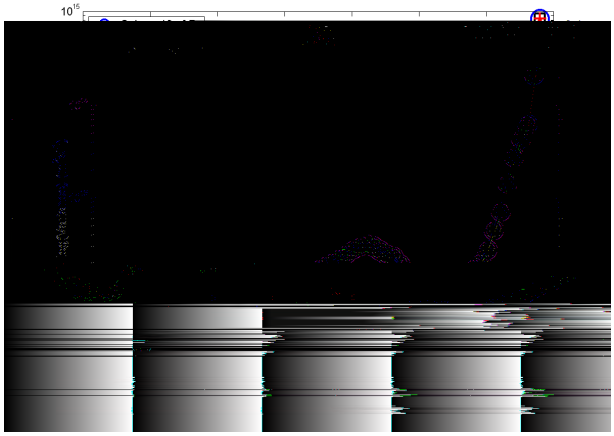

Figure 4.6: Plot in log-scale showing the three columns of R calculated using the `delta' test with the ECMWF Spatial Transform. The blue line corresponds to column 10. Column 50 is plotted in black and column 80 in red.

4.3 Observing Separability of Correlation Functions

In order to investigate the separability of the vertical correlation functions resulting from the spatial transforms, the correlations between the di erent vertical levels are observed for a speci c level. To achieve this requires the calculation of the implied vertical correlation matrix C^{imp} 2 $R^{IJ \times IJ}$. This matrix is related to the B^{imp} matrix through the matrix 2 $R^{IJ \times IJ}$, a diagonal matrix with the model level standard deviations along the diagonal and zeros elsewhere. C^{imp} is found using

$$C^{imp} = {}^{-1} B^{imp} {}^{-1}. \tag{4.1}$$

In order to observe the correlations using the R matrix calculated using the spatial transforms in Chapter 3, the correlation matrix C^R 2 $R^{I \times J}$ is calculated. This uses the matrix de ned in equation (3.41) of Section 3.3 and the matrix as described in Section 3.4. To observe the correlations with respect to a particular level l at horizontal grid-point n the matrix is constructed such that

$$_{(i;j)} = \begin{cases} 1 & \text{if } i = l; j = n; \\ 0 & \text{otherwise.} \end{cases} \tag{4.2}$$

This matrix is then pre-multiplied by and the R matrix is constructed using the vertical and horizontal transforms as described in Sections 3.2 or 3.3. R is then pre-multiplied by to produce the matrix C^R 2 $R^{I \times J}$. This gives a representation of the correlation at each grid-point with respect to the point $(l; n)$. This can be written in the form

$$C^R = R \quad ; \tag{4.3}$$

where R is an operator signifying the construction of the matrix R.

By altering the value of l for which C^R is calculated the horizontal and vertical length scales can be observed. Figures 4.7 - 4.9 show a small selection of the resulting correlations showing how the correlations change as the vertical level increases, using the Met O ce Spatial Transform. The results are shown for horizontal grid-point 50 in order to observe the resulting shapes more clearly. The choice of horizontal grid-point is irrelevant due to the homogeneity assumption used in calculating B^{imp} and R.

The correlations in this section are plotted with respect to log(Pressure), rather than vertical level. This is to ensure that the length scales can be observed more accurately, as the vertical levels are not uniformly spaced (see Table 1). As a result of plotting against log(Pressure), the lower vertical levels are shown at the top of the plots and the higher vertical levels at the bottom of the plot.

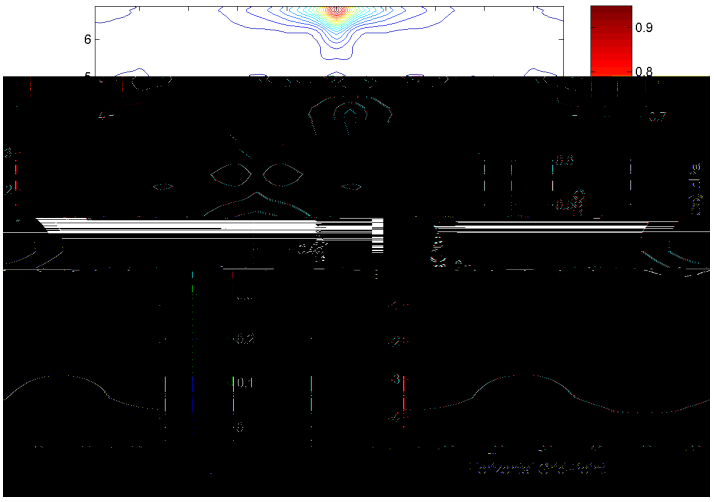

Figure 4.7: Correlations between vertical levels for level 11 from the Met O ce C^R matrix.

Figure 4.7 shows the correlation with vertical level 11 at horizontal grid-point 50. The correlation at the point observed is equal to one as expected. The correlation bands are narrow in this case, with the values dropping toward zero sharply in the horizontal direction. The vertical and horizontal length scales are small, which gives an indication that the functions are non-separable. Separable correlation functions imply that the horizontal and vertical length scales need not be equal.

This similarity between the length scales also demonstrates that the correlations are isotropic in grid-point space. The correlations are not circular, which could be a result of the coarse resolution being used. There is also a small amount of correlation in levels higher than that level 11 (in the log(Pressure) range (0 1) and (3 4) hPa in Figure

4.7). This is an example of the unphysical behaviour seen in the top ve levels in Figure 4.1 and will be a feature of all the plots shown in this section.

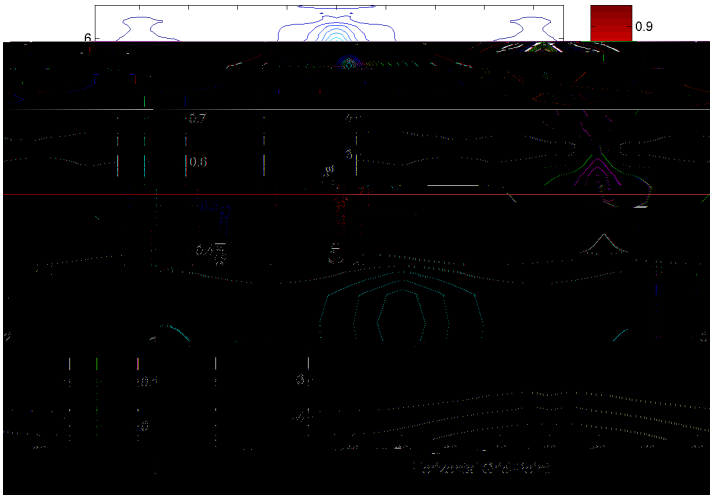

Figure 4.8: Correlations between vertical levels for level 41 from the Met O ce C^R matrix.

Figure 4.8 is a representation of the correlations in the mid-vertical levels, with

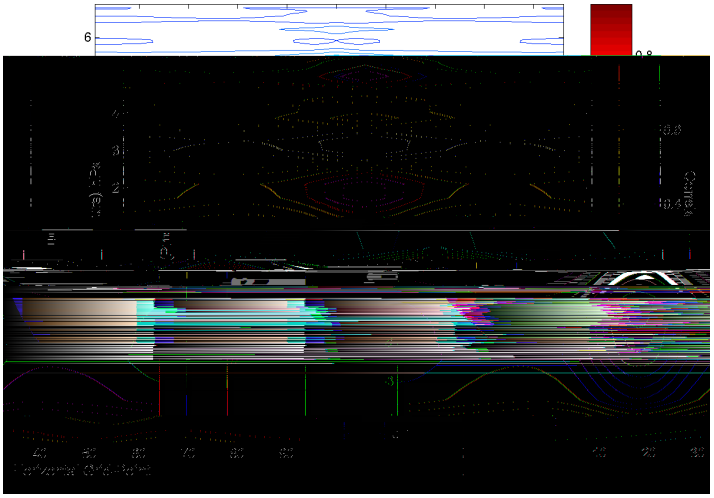

Figure 4.9: Correlations between vertical levels for level 66 from the Met O ce C^R matrix.

The results presented give a good indication of the non-separability of the correlation functions for the Met O ce spatial transform. By comparing these results with the resulting correlations from the ECMWF spatial transform a comparison between the two can be made. Figures 4.10 - 4.12 show the correlation functions for the same vertical levels as presented for the Met O ce, but using the ECMWF spatial transform.

Figure 4.10 show the resulting correlations for vertical level 11. An immediate comparison can be made with Figure 4.7 in that the vertical and horizontal length scales

porizontal tength tcales-395(textedicg)-3175(m28(tuc28(t)-395(tfurhe)r-395(tacross-395(the)-3165poriz

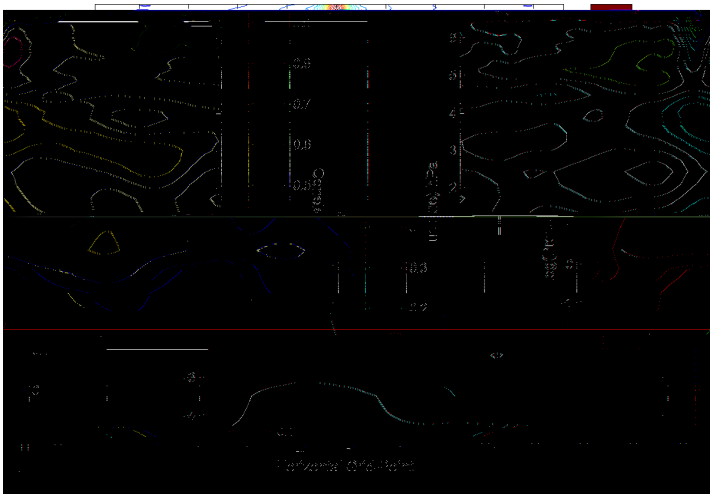

Figure 4.10: Correlations between vertical levels for level 11 from the ECMWF C^R matrix.

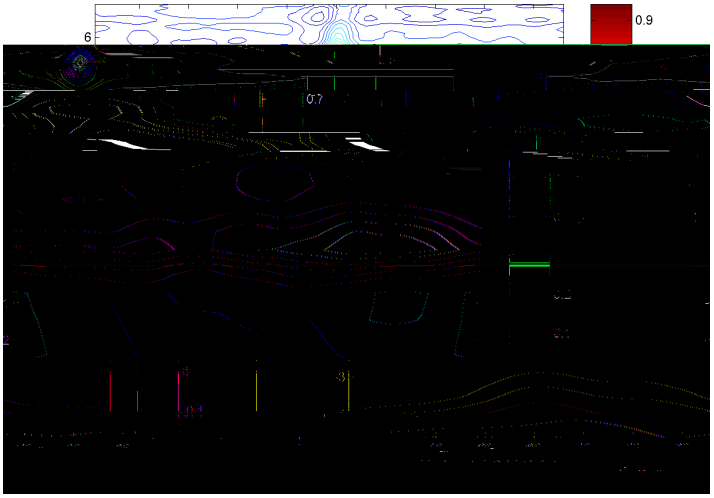

Figure 4.11: Correlations between vertical levels for level 41 from the ECMWF C^R matrix.

the Met O ce example, which could be a result of the di erent way in which the spatial transform is constructed.

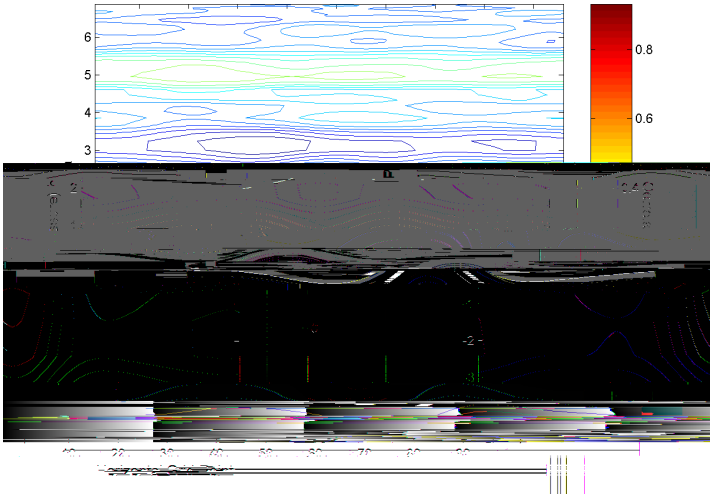

Figure 4.12: Correlations between vertical levels for level 66 from the ECMWF C^R matrix.

For the top ve levels, of which level 66 is given as an example in Figure 4.12, the horizontal length scale extends across the whole of the horizontal direction with higher than expected correlations. If compared to Figure 4.9, the correlations with the lower vertical levels are higher, but still extend along all of the horizontal direction. This indicates that the horizontal length scales for the higher vertical levels are larger than

the horizontal and vertical directions. Bartello and Mitchell (1992) and Lonnberg and Hollingsworth (1986) also provides results showing the increase in horizontal length scales as the vertical level increases. The sharpening of vertical levels with increased height was shown by Derber and Bouttier (1999), which further highlights the non-separability discussed here.

4.4 Vertical Mode Analysis

The ECMWF spatial transform relies upon the storage of much more information than in the Met O ce spatial transform, but as has been seen in Section 4.2 the results produced as a result of the `delta' test are the same. In order to understand the reasoning behind the ECMWF storing this extra data it is possible to examine the dominant vertical modes, which are the eigenvectors found in the vertical transform stage (see Sections 3.2.2 and 3.3.2). These are the modes that are associated with the eigenvalues of greatest magnitude. Vertical mode analysis has also been performed by Ingleby (2001) for a variety of control variables.

Since the ECMWF stores J pseudo-correlation matrices, one for each wavenumber, for each of the I vertical modes there are J ECMWF vertical modes compared to one in the Met O ce. In order to investigate the di erence between the two operational centres a selection of these are plotted. Since the vertical modes are in spectral space, they consist of a real and imaginary part, both of which need to be investigated.

Figure 4.13 is a plot of four of the vertical modes associated with the rst ECMWF vertical mode, representing horizontal wavenumbers 1, 15, 30 and 48. The real part of the mode is plotted in the green dashed line, while the imaginary part is the red dotted line. The solid blue line is the rst Met O ce vertical mode. The plot corresponding to wavenumber 1 has no imaginary part as this mode is real-valued despite being in spectral space. Due to the complex conjugacy property discussed in Section 3.1.4, only the rst half of the spectrum needs considering.

The plot shows that for the lower and higher horizontal wavenumbers the components of the eigenvectors are largest (in absolute value) for the higher vertical levels and

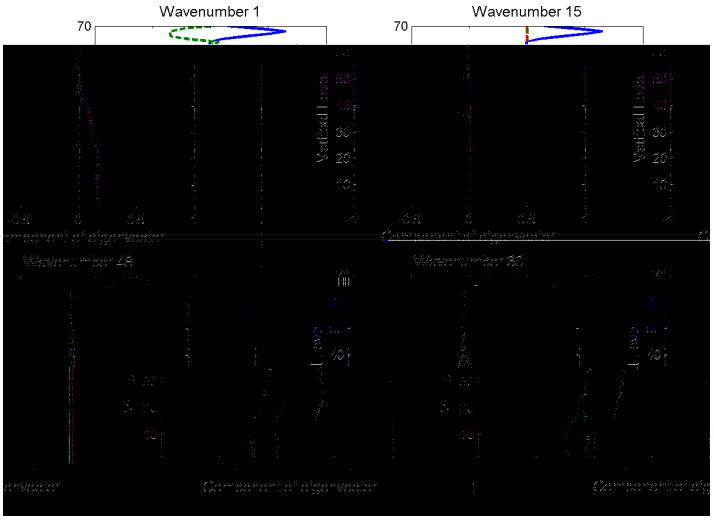

Figure 4.13: Plots of 4 horizontal wavenumbers in the ECMWF vertical transform corresponding to the first vertical mode. Green dashed line is the real part of the mode, red dash-dotted line is the imaginary part. The Met Office mode is plotted (solid blue line) for comparison.

smaller for the lower levels. The horizontal wavenumbers toward the middle of the range (wavenumbers 30 and 48 in Figure 4.13) have the largest magnitude of components for the lower vertical levels and almost no magnitude in the higher levels.

This feature of the modes would indicate that the ECMWF vertical transform ensures that all the vertical levels have some contribution made for the leading modes. The Met Office leading modes have almost no contribution to the lower levels, resulting in the majority of the power being present in the higher vertical levels (since the leading modes are associated with eigenvalues of greatest magnitude).

Performing the same analysis for other leading modes justifies this result, with the third dominant mode shown in Figure 4.14. The mid-range modes (mode 35 for example, not shown) have a contribution that is more uniformly distributed throughout the vertical levels in both operational centres. For the vertical modes corresponding to the eigenvectors of smallest magnitude (e.g. mode 67, not shown) the contribution to the

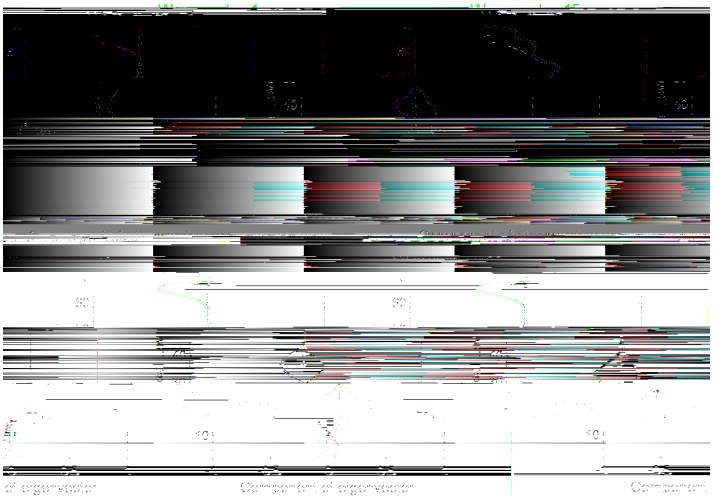

Figure 4.14: Plots of 4 horizontal wavenumbers in the ECMWF vertical transform corresponding to the third dominant mode. Green dashed line is the real part of the

meaning. In order to investigate the way in which this impacts on the resulting B^{imp} matrix, a comparison between the results obtained from the Met O ce R matrix is made using both of the outer products described.

4.5.1 Applying the `delta' test

for the rst vertical mode the majority of the power in the mode is attributed to the lower wavenumbers, decreasing slowly as the wavenumber increases. The results from the two outer products are similar, with no obvious di erence between the two. This is common to the other low vertical modes (not plotted), indicating that the more physical representation of the vertical covariance matrix has little impact upon the lower vertical modes.

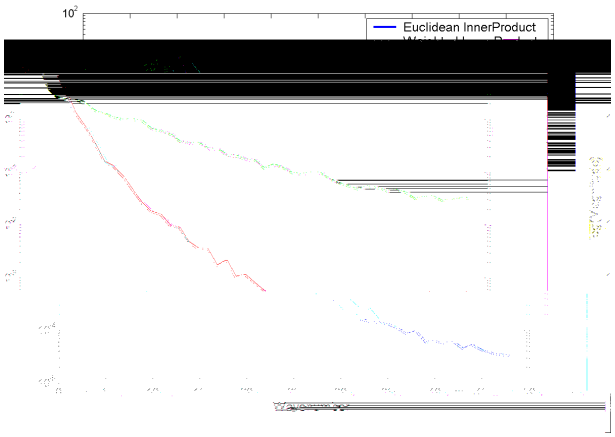

Figure 4.17: The Power Spectrum for vertical mode 67 calculated using the Euclidean outer product (blue line) and the mass-weighted outer product (dashed red line).

This similarity between the two outer products continues for the majority of the vertical modes, with little di erence visible (some exceptions occur during the mid-range vertical modes but this is not signi cant to look at more deeply here). When the highest vertical modes are reached (65 upwards) a marked di erence between the two outer products begins to appear for the wavenumbers. Figure 4.17 shows the power spectrum for vertical mode 67.

The power spectrum in this case is much atter for the mass-weighted outer product, showing that the power has been distributed more evenly about the wavenumbers. This result shows that the representation of the higher vertical modes is more dynamically reasonable. It can also be used to show that the correlation functions are more separable than for the Euclidean outer product de nition. This links to Section 4.3 and a similar

test for the mass-weighted outer product could be performed.

This investigation has gone some way toward answering question 3 at the beginning of the Chapter. By implementing the mass-weighted outer product the power spectra for the higher vertical modes have become much atter and have therefore produced a more physical representation of the vertical transforms. This section has also linked in which the discussion on the separability of correlation function.

4.6 Summary

In this Chapter a number of investigations were performed in order to answer three questions relating to the implied B matrix. These investigations involve the use of streamfunction sample data provided by the Met O ce.

Examining the grid-point variances shows an increase in the variance through the vertical levels, with a sharp increase at the top ve levels. The results of the `delta' test, as described in Section 3.4, show that the only di erence between the Met O ce and ECMWF spatial transforms can be attributed to round-o error. The homogeneity assumption is also validated for both the Met O ce and ECMWF spatial transforms.

By observing the correlations between vertical levels for a selection of vertical levels insight was gleaned concerning the non-separability of the correlation functions. It was seen that both the vertical and horizontal length scales increase as vertical level increases and are isotropic in the case of the Met O ce transform. The change in length scale is more pronounced in the ECMWF transform than that of the Met O ce.

Vertical mode analysis showed that for the ECMWF, the speci cation of a set of eigenvalues and eigenvectors for each horizontal wavenumber has an e ect on the leading vertical modes. In the Met O ce the leading modes have a greater magnitude for the higher vertical levels. The ECMWF has magnitude that is spread throughout the vertical levels, depending upon the wavenumber observed.

The e ect on implementing the mass-weighted outer product into the Met O ce vertical transforms was highlighted through analysis of the power spectrum. By comparing the power spectrum with that obtained through the standard Euclidean outer

product, it was shown that for higher vertical modes the power spectrum is signi cantly atter for the mass-weighted outer product. This is a bene t of incorporating a more dynamical and physical representation into the vertical transform.

Observation of the power spectra also highlighted a possibility of the correlation functions being more separable in the mass-weighted representation, which would describe a more realistic correlation structure.

spatial transform. since the horizontal and vertical length scales were of similar size. Some discrepancy occurs but is possibly due to the coarse resolution on which the data is prescribed. The isotropy is not so clear in the ECMWF spatial transform, which was an unexpected result as the assumption forms a key part of the spatial transform.

Vertical mode analysis provided some justi cation for the increased storage seen in the ECMWF vertical transforms. The Met O ce leading vertical modes have almost the entire magnitude of the mode attributed to the highest vertical levels (also seen in Ingleby (2001)). By plotting a selection of modes corresponding to di erent wavenumber for the ECMWF leading modes, it was seen that the distribution of magnitude is spread throughout the vertical levels. This more uniform distribution of vertical mode contribution is a bene t of the ECMWF additional storage.

Incorporation of a mass-weighted outer product into the vertical transform stage of the Met O ce spatial transform produced some surprising results. The initial results obtained testing the resulting implied B matrix suggest that there is no di erence between the twoorresp2.90197.icatdt28(w)2d(b)-2n

sented here only represent a small proportion of those that could be performed in order to answer the three main questions posed in Chapter 1. There is also much potential in extending the investigations already performed in order to answer these questions more thoroughly.

The assumptions of homogeneity and isotropy play an important role in the implied B matrix resulting from the spatial transforms. Further analysis of these assumptions is achievable, both in terms of the sample data and resulting covariances. This could be achieved through the use of single observation tests, which would highlight the analysis increment around the observation. It would be expected that the observation information would spread to nearby point in a symmetric manner if the assumptions of isotropy hold.

The potential of the mass-weighted outer product has not been fully explored, particularly in relation to the ECMWF spatial transform. Analysis of the e ect on the resulting vertical transforms would be worth performing, especially considering the differences between the way in which they are constructed in the two operational centres. The vertical mode analysis presented here could then also be applied to the mass-weighted outer product, where it would be expected that the leading modes would be di erent to those seen in this project.

The higher vertical levels were the cause of some strange results observed throughout this project, some of which cannot be explained using physical reasoning. By truncating the eld to remove these vertical levels, a more physical picture of the atmosphere could be observed. This would enable the analysis to be performed without the strange behaviour becoming a factor in the results.

In presenting a simple two-dimensional eld upon which the sample data was collected, limitations were placed upon the experiments that could be performed. A logical extension would be to allow for variations in latitude for the sample data, which would

horizontal transform however, as it would involve spherical harmonics rather than the DFT used here.

Finally, this project involved the use of only a single control variable, allowing only univariate covariances to be examined. By introducing more control variables into the project, a multivariate analysis could be performed. This would allow the relation between di erent variables to be explored, which was not possible in this project.

Appendix

Model Level Height/Pressure Specifications

In table 1 on the following page the model level heights and equivalent pressures are given.

Level	Height (m)	Equivalent pressure (hPa)	Level	Height (m)	Equivalent pressure (hPa)
1	10.000	983.51	36	8877.3	308.87
2	36.664	979.75	37	9371.4	287.37
3	76.664	974.51	38	9879.6	266.54
4	130.00	967.81	39	10402	246.43
5	196.66	959.66	40	10940	227.07
6	276.66	950.10	41	11494	208.49
7	370.00	939.17	42	12065	190.69
8	476.66	926.90	43	12654	173.68
9	596.66	913.33	44	13265	157.45
10	730.00	898.53	45	13898	141.98
11	876.66	882.55	46	14558	127.27
12	1036.7	865.45	47	15249	113.31
13	1210.0	847.30	48	15975	100.13
14	1396.7	828.17	49	16743	87.77
15	1596.7	808.12	50	17559	76.28
16	1810.0	787.24	51	18432	65.70
17	2036.7	765.59	52	19372	56.01
18	2276.7	743.26	53	20392	47.17
19	2530.0	720.31	54	21505	39.17
20	2796.7	696.81	55	22728	31.99
21	3076.7	672.86	56	24079	25.62
22	3370.0	648.51	57	25580	20.06
23	3676.7	623.86	58	27256	15.31
24	3996.7	598.96	59	29135	11.34
25	4330.0	573.91	60	31250	8.12
26	4676.7	548.78	61	33637	5.61

Bibliography

E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK User's Guide*. SIAM, third edition, 1999.

H. Anton and C. Rorres. *Elementary Linear Algebra- Applications Version*. J.Wiley and Sons Inc., eighth edition, 2000.

R.N. Bannister. A review of forecast error covariance statistics in atmospheric variational data assimilation. I: Characteristics and measurements of forecast error covariances. *Q.J.R. Meteorol. Soc.*, 134:1951{1970, 2008a.

R.N. Bannister. A review of forecast error covariance statistics in atmospheric variational data assimilation. II: Modelling the forecast error covariance statistics. *Q.J.R. Meteorol. Soc.*, 134:1971{1996, 2008b.

R.G. Barry and R.J. Chorley. *Atmosphere, Weather and Climate*. Routledge, ninth edition, 2010.

P. Bartello and H.L. Mitchell. A continuous three-dimensional model of short-range forecast error covariances. *Tellus*, 44A:217{235, 1992.

G.M. Baxter. *4D-Var for high resolution, nested models with a range of scales*. PhD thesis, University of Reading, 2009.

F. Bouttier. A dynamical estimation of forecast error covariances in an assimilation system. *Mon. Weather Review*, 122:2376{2390, 1994.

F. Bouttier and P. Courtier. Data assimilation concepts and methods. ECMWF Meteorological Training Course Lecture Series, 1999. URL www.ECMWF.int. Last Accessed 19/08/2010.

M. Buehner. Ensemble-derived stationary and ow-dependent background-error covariances: Evaluation in a quasi-operational NWP setting. *Q.J.R. Meteorol. Soc.*, 131: 1013{1043, 2005.

P. Courtier. Dual formulation of four-dimensional variational assimilation. *Q.J.R. Meteorol. Soc.*, 123:2449{2461, 1997.

P. Courtier, J.N. Thepaut, and A. Hollingsworth. A strategy for operational implementation of 4D-Var, using an incremental approach. *Q.J.R. Meteorol. Soc.*, 120: 1367{1387, 1994.

P. Courtier, E. Andersson, W. Heckley, J. Pailleux, D. Vasiljevic, M. Hamrud, A. Hollingsworth, F. Rabier, and M. Fisher. The ECMWF implementation of three-dimensional variational assimilation (3D-Var). I: Formulation. *Q.J.R. Meteorol. Soc.*, 124:1783{1807, 1998.

R. Daley. *Atmospheric Data Analysis*. Cambridge University Press, 1991.

J. Derber and F. Bouttier. A reformulation of the background error covariance in the ECMWF global data assimilation system. *Tellus*, 51A:195{221, 1999.

G. Desroziers. A coordinate change for data assimilation in spherical geometry of frontal structures. *Mon. Weather Rev.*, 125:3030{3038, 1997.

M. Fisher. Background error covariance modelling. In *ECMWF Seminar on Recent Developments in Data Assimilation for Atmosphere and Ocean*, pages 45{64, ECMWF: Reading, UK, September 2003.

M. Fisher and P. Courtier. Estimating the covariance matrices of analysis and forecast error in variational data assimilation. Technical Memorandum no. 220, ECMWF

Research Dept., Shin eld Park, Reading, RG2 9AX, UK, 1995. Available from

N.B. Ingleby. The statistical structure of forecast errors and its representation in the met. o ce global 3-d variational data assimilation scheme. *Q.J.R. Meteorol. Soc.*, 127:209{231, 2001.

A. Iserles. *A First Course in the Numerical Analysis of Di erential Equations*. Cambridge University Press, second edition, 2009.

E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2003.

D. Kincaid and W. Cheny. *Numerical Analysis: Mathematics of Scienti c Computing*. Brooks/Cole, third edition, 2002.

J.M. Lewis, S. Lakshmivarahan, and S.K. Dhall. *Dynamic Data Assimilation: A Least Squares Approach*. Cambridge University Press, 2006.

P. Lonnberg and A. Hollingsworth. The statistical structure of short-range forecast errors as determined from radiosonde data part II: The covariance of height and wind errors. *Tellus*, 38A:137{161, 1986.

A.C. Lorenc. Analysis methods for numerical weather prediction. *Q.J.R. Meteorol. Soc.*, 112:1177{1194, 1986.

A.C. Lorenc. Modelling of error covariances by 4D-Var data assimilation. *Q.J.R. Meteorol. Soc.*, 129:3167{3182, 2003a.

A.C. Lorenc. The potential of the ensemble Kalman lter for NWP- a comparison with 4D-Var. *Q.J.R. Meteorol. Soc.*, 129:3183{3203, 2003b.

A.C. Lorenc, S.P. Ballard, R.S. Bell, N.B. Ingleby, P.L.F. Andrews, D.M. Barker, J.R. Bray, A.M. Clayton, T. Dalby, D. Li, T.J. Payne, and F.W. Saunders. The Met O ce global three-dimensional variational data assimilation scheme. *Q.J.R. Meteorol. Soc.*, 126:2991{3012, 2000.

T.L. Otte, N.L. Seaman, and D.R. Stau er. A heuristic study on the importance of anisotropic error distributions in data assimilation. *Mon. Weather Review*, 129:766{ 783, 2001.