School of Mathematical
and Physical Sciences

University of Reading

Department of Mathematics and Statistics

Preprint MPS-2012-18

3 September 2012

Variational data assimilation for very large environmental problems

by

A.S. Lawless

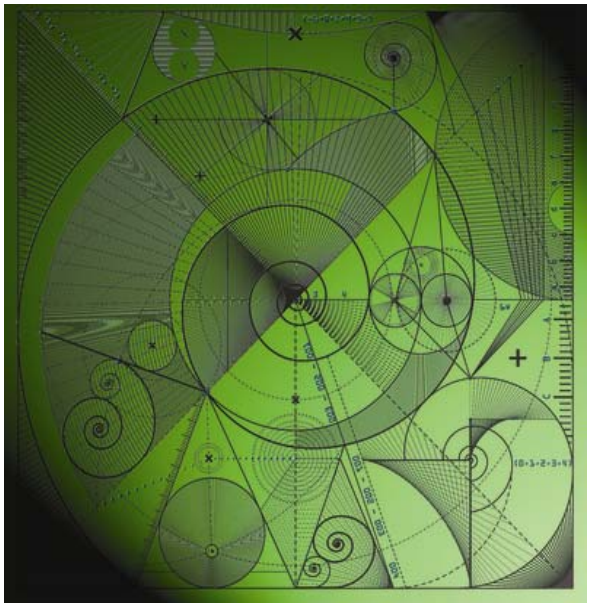

Variational data assimilation for very large environmental problems

A.S. Lawless*

School of Mathematical and Physical Sciences, University of Reading, PO Box 220, Whiteknights, Reading, RG6 6AX
a.s.lawless@reading.ac.uk

Abstract

Variational data assimilation is commonly used in environmental forecasting to estimate the current state of the system from a model forecast and observational data. The assimilation problem can be written simply in the form of a nonlinear least squares optimization problem. However the practical solution of the problem in large systems requires many careful choices to be made in the implementation. In this article we present the theory of variational data assimilation and then discuss in detail how it is implemented in practice. Current solutions and open questions are discussed.

KEYWORDS: 3D-Var, 4D-Var, adjoint model, background errors, error covariance, incremental formulation, nested models, observation errors, optimization, reduced order models, tangent linear model, weak-constraint.

1 Introduction

Data assimilation is the process of combining a numerical model forecast with observational data in order to estimate the current state of a dynamical system. It has been an essential part of numerical weather prediction (NWP) since its beginnings in the 1940s, when it was recognized that errors in the initial model state could rapidly lead to large errors in the forecast. Early data assimilation schemes were based on a simple interpolation between the observations and the model state, with later schemes also taking account of the statistics of the errors in the data. Such schemes included smoothing splines, successive correction, optimal interpolation and analysis correction [68], [71]. The possible use of methods based on variational calculus was proposed by Sasaki [86], [87]

*The author is supported in part by the U.K. Natural Environment Research Council, through the National Centre for Earth Observation.

in the late 1950s and 1960s, but at the time a practical implementation was not possible. A real breakthrough in the application of variational schemes to NWP came in the late 1980s with a series of papers demonstrating how the problem could be solved using techniques from the theory of optimal control, in particular the use of adjoint equations to calculate the gradient of an objective function, or cost function [63], [90]. This led to a series of papers in which the feasability of variational data assimilation was studied on a series of simplified atmospheric models [91], [81], [18], [77] (these experiments usually only included the large-scale atmospheric dynamics and not the subgrid-scale processes of full weather prediction models).

Despite the encouraging results of these experiments variational data assimilation remained impractical for operational use due to the high computational cost. The introduction of the incremental method of variational assimilation in 1994 [19], together with increasing computing power, opened up the possibility of an affordable implementation for operational weather prediction. Over the following decade many weather forecasting centres began to develop variational data assimilation for operational use [82], [70], [83], [31], [32], [48]. At the same time variational data assimilation began to be applied to other applications, such as ocean forecasting [99], [95] and atmospheric chemistry [28].

A common feature of many of these applications is that the size of the state variable being estimated is extremely large. Current numerical weather prediction models may require the initialization of the order of 10^8 variables in order to make a forecast. As computing power increases the spatial resolution of the models tends to increase and hence so does the number of variables being represented. Furthermore the real-time nature of environmental forecasting requires that the data assimilation problem be solved quickly. These two factors imply that when implementing variational data assimilation schemes in practice compromises must be made. Hence it is important to design the algorithms carefully to ensure that as accurate a solution as possible is obtained within the time available. Ideally such design should also include knowledge of the physics of the problem, so that the final solution is physically realistic. In the remainder of this article we will discuss some of the different choices that arise in the implementation of variational data assimilation for very large systems and the practical approaches that have been developed. First we briefly present the mathematical theory of variational data assimilation.

2 Theory of variational data assimilation

We consider a discrete nonlinear dynamical system given by the equation

$$\mathbf{x}_{i+1} = \mathcal{M}_i(\mathbf{x}_i), \tag{1}$$

where $\mathbf{x}_i \in \mathbb{R}^n$ is the state vector at time t_i and $\mathcal{M}_i$ is the nonlinear model operator that propagates the state at time t_i to time t_{i+1} for $i = 0, 1, \ldots, N-1$. We assume that we have imperfect observations $\mathbf{y}_i \in \mathbb{R}^{p_i}$ at times $t_i, i = 0, \ldots, N$

that are related to the model state through the equation

$$\mathbf{y}_i = \mathcal{H}_i(\mathbf{x}_i) + \boldsymbol{\eta}_i, \tag{2}$$

where $\mathcal{H}_i : \mathbb{R}^{p_i} \to \mathbb{R}^n$ is known as the observation operator and maps the state vector to observation space. The observation errors $\boldsymbol{\eta}_i$ are usually assumed to be unbiased, serially uncorrelated, Gaussian errors with known covariance matrices $\mathbf{R}_i$. For the numerical weather prediction problem the vector $\mathbf{x}_i$ would contain several meteorological variables, such as pressure, temperature and the three-dimensional wind, at each grid point of the model domain. The observation operator $\mathcal{H}_i$ may just be a simple interpolation in space, if the state variable is observed directly. However, it could be a much more complicated nonlinear function of the state. For example, for a satellite radiance measurement the observation operator can include a complex radiative transfer model.

We assume that at the initial time t_0 we have an *a priori* estimate of the state, usually referred to as a *background* field, that we denote $\mathbf{x}^b$. This background field is assumed to have unbiased, Gaussian errors with known covariance matrix $\mathbf{B}$. In practice the background field is usually a short-term forecast of the state from a previous assimilation cycle. The problem of four-dimensional variational data assimilation (4D-Var) is then to find the initial state that minimizes the weighted least squares distance to this background while minimizing the weighted least squares distance of the model trajectory to the observations over the time interval $[t_0, t_N]$ [1]. Mathematically we can formulate this as an optimization problem:
Find the state $\mathbf{x}_0^a$ at time t_0 that minimizes the function

$$\mathcal{J}(\mathbf{x}_0) = \frac{1}{2}(\mathbf{x}_0 - \mathbf{x}^b)^{\mathrm{T}}\mathbf{B}^{-1}(\mathbf{x}_0 - \mathbf{x}^b) + \frac{1}{2}\sum_{i=0}^{N}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i)^{\mathrm{T}}\mathbf{R}_i^{-1}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i) \tag{3}$$

subject to the states $\mathbf{x}_i$ satisfying the nonlinear dynamical system (1). In the case where $N = 0$ there is no model evolution and the scheme is referred to as three-dimensional variational data assimilation (3D-Var). The solution $\mathbf{x}_0^a$ is commonly referred to as the *analysis*. In environmental data assimilation the function $\mathcal{J}(\mathbf{x}_0)$ is usually called the *cost function*, but the terms *objective function* and *penalty function* are often used in other fields.

The minimization problem given by equation (3) can be interpreted in a statistical or deterministic sense. From Bayes' theorem it can be shown that $\mathbf{x}_0^a$ gives the maximimum likelihood estimate of the state under the assumptions given [68]. Alternatively the term measuring the fit to the background state can be thought of as a form of Tikhonov regularization in fitting the observations [52], [21]. Each of these interpretations is able to provide different insights into the practical formulation of the problem.

The cost function can be minimized using iterative numerical methods, such as conjugate gradient or quasi-Newton methods. The use of these methods in

[1] The scheme is referred to as four-dimensional since we usually fit three spatial dimensions in time, with time being the fourth dimension.

data assimilation is discussed in more detail in section 3.4. On each iteration of such methods the value of the cost function and its gradient at the current iterate must be calculated. The gradient of (3) with respect to the initial state $\mathbf{x}_0$ can be found by first solving the discrete adjoint equations

$$\boldsymbol{\lambda}_i = \mathbf{M}_i^T \boldsymbol{\lambda}_{i+1} - \mathbf{H}_i^T \mathbf{R}^{-1}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i) \tag{4}$$

where $\boldsymbol{\lambda}_i$ are the adjoint variables, with $\boldsymbol{\lambda}_{N+1} = 0$, and $\mathbf{H}_i$ and $\mathbf{M}_i$ are the Jacobians of the nonlinear operators $\mathcal{H}_i$ and $\mathcal{M}_i$ with respect to the state variable $\mathbf{x}_i$. In the data assimilation literature these Jacobians are referred to as the *tangent linear operator* and the *tangent linear model* (TLM). The gradient of the cost function with respect to the initial state is then given by

$$\nabla \mathcal{J}(\mathbf{x}_0) = -\boldsymbol{\lambda}_0 + \mathbf{B}^{-1}(\mathbf{x}_0 - \mathbf{x}^b), \tag{5}$$

where the operators $\mathbf{H}_i^T$ and $\mathbf{M}_i^T$ are the adjoints of the observation operator and the nonlinear model. We note that these adjoints are usually taken with respect to the Euclidean inner product and therefore the adjoint is equivalent to the matrix transpose of the Jacobians. Other inner products are only necessary where a physical interpretation of the adjoint variables is required. Each iteration of a numerical optimization method therefore requires one run of the forward model (1) to calculate the value of the cost function and one run of the adjoint model (4) to calculate the gradient. This makes 4D-Var very expensive from a computational point of view.

2.1 Incremental variational data assimilation

The possibility of implementing variational data assimilation in an operational setting came with the proposal of incremental variational data assimilation [19]. In this formulation the solution to the nonlinear miminimization problem (3) is approximated by a sequence of minimizations of linear quadratic cost functions. We define $\mathbf{x}_0^{(k)}$ to be the k^{th} estimate to the solution and linearize the cost function (3) around the model trajectory forecast from this estimate. The next estimate is then defined by

$$\mathbf{x}_0^{(k+1)} = \mathbf{x}_0^{(k)} + \delta\mathbf{x}_0^{(k)}, \tag{6}$$

where the perturbation $\delta\mathbf{x}_0^{(k)} \in \mathbb{R}^n$ is a solution of the linearised cost function

$$\begin{aligned} \tilde{\mathcal{J}}^{(k)}(\delta\mathbf{x}_0^{(k)}) &= \frac{1}{2}(\delta\mathbf{x}_0^{(k)} - [\mathbf{x}^b - \mathbf{x}_0{}^{(k)}])^{\mathrm{T}}\mathbf{B}^{-1}(\delta\mathbf{x}_0^{(k)} - [\mathbf{x}^b - \mathbf{x}_0{}^{(k)}]) \\ &+ \frac{1}{2}\sum_{i=0}^{N}(\mathbf{H}_i\delta\mathbf{x}_i^{(k)} - \mathbf{d}_i^{(k)})^{\mathrm{T}}\mathbf{R}_i^{-1}(\mathbf{H}_i\delta\mathbf{x}_i^{(k)} - \mathbf{d}_i^{(k)}). \end{aligned} \tag{7}$$

Here $\mathbf{d}_i^{(k)} = \mathbf{y}_i - \mathcal{H}_i(\mathbf{x}_i^{(k)})$, where $\mathbf{x}_i^{(k)}$ is the nonlinear trajectory calculated from the current estimate at the initial time using the nonlinear model equation (1). The perturbation $\delta\mathbf{x}_i$ satisfies the linear dynamical equation

$$\delta\mathbf{x}_{i+1} = \mathbf{M}_i\delta\mathbf{x}_i. \tag{8}$$

The linearized observation operator $\mathbf{H}_i$ and the tangent linear model operator $\mathbf{M}_i$ are evaluated at the current estimate of the nonlinear trajectory, usually called the linearization state. The minimization (7) is referred to as the *inner loop*, while the update of the nonlinear model trajectory $\mathbf{x}_i^{(k)}$ is the *outer loop*. On each iteration of the inner loop the TLM is integrated to calculate the evolution of the perturbation, in order to calculate the cost function (7), and the adjoint model is integrated to provide the gradient. A major advantage of the incremental approach is that the inner loop minimization problem may be solved in a smaller dimensional space than the outer loop step, for example at a lower spatial resolution. In this way the TLM and adjoint model need only be run at the lower resolution on each inner loop iteration, while the linearization trajectory from the nonlinear model is still calculated at the higher resolution on each outer loop. The computational savings made by implementing the inner loop in this way made incremental 4D-Var feasible for operational weather and ocean forecasting.

The incremental method was later shown to be equivalent to an inexact Gauss-Newton method applied to the original nonlinear cost function (3) [58]. The outer loop iterations can be shown to be locally convergent under certain conditions, provided that the inner loop minimization is solved to sufficient accuracy [57], [34]. In practice, however, very few outer loop steps are performed. For example, the Met Office perform only one, while the European Centre for Medium-range Weather Forecasts (ECMWF) perform three [83], [29]. As for the fully nonlinear problem the incremental method can be run as 3D-Var (no model evolution) or 4D-Var (including the model evolution). An alternative formulation that is often implemented is known as 3D-FGAT (First Guess at Appropriate Time). This includes the nonlinear model evolution in the calculation of the vectors $\mathbf{d}_i$, but no evolution is included for the perturbation and the TLM operator $\mathbf{M}_i$ in equation (8) is replaced by the identity. This ensures that the observations are compared with the nonlinear trajectory at the correct time, but approximates the perturbation in such a way that no TLM or adjoint model is needed. In this way some of the benefit of 4D-Var can be achieved without too much extra compuational cost [73], [56].

Having presented the basic theory of variational data assimilation we now examine some of the issues that arise in its practical implementation. For the very large systems found in environmental modelling it is not always possible to apply the theory in an intuitive way. Many choices must be made in order to set up and solve the assimilation problem efficiently and compromises must often be made. It is the attention to detail in these choices that can determine the success or otherwise of the data assimilation scheme.

3 Practical implementation

3.1 Model development

The development of a 4D-Var scheme for the large models used in operational weather and ocean forecasting is often a huge undertaking. In most cases the nonlinear model code already exists and has been developed over many years. These models are very large pieces of software, with maybe close to one million lines of code. In order to develop an incremental 4D-Var scheme the code for the TLM and adjoint model must first be written. The development of a TLM code and adjoint model code from the source code of a nonlinear model is a fairly automatic procedure. The correct code for the TLM can be found from a linearization of each statement of the nonlinear model source code, based on treating the nonlinear model as a series of arithmetic operations and applying the chain rule. The adjoint model is then found by a line-by-line transpose of the TLM source code in reverse order. This method is known as automatic differentiation. We do not go into details of its application here, but refer the reader to several good introductions in the literature [18], [8], [85], [33]. The automatic nature of this procedure has led to many software tools being developed that will produce a TLM and adjoint model code from a nonlinear mode source code. These automatic differentiation tools, or automatic adjoint compilers, are now available commercially for many different programming languages. [2]

In practice the TLM and adjoint models of many large environmental models have been developed by hand, rather than using the automatic compilers. There are several reasons for this. The first is that in many cases of operational weather and ocean forecasting the complexity of the already exisiting nonlinear model codes was such that simple application of the automatic compilers was not possible. In many cases, particularly for large codes developed by many people, it is necessary to tidy the nonlinear model codes to make them suitable for use with the automatic compilers. Many centres felt that the effort to do this would have been greater than coding the TLM and adjoint model by hand.

The second reason for developing the TLM and adjoint codes by hand arises from the nature of the incremental approach to variational data assimilation. Since the TLM and adjoint are run at a lower resolution in the inner loop, the TLM is already an approximate linearization of the nonlinear model used in the outer loop. It is therefore justifiable to make further simplifications in the TLM, in order to reduce the computational cost. As long as the adjoint model is derived from the approximate TLM, then the inner loop minimization will contain the correct gradient information for convergence. In coding the models by hand it is easier to make such simplifications based on physical arguments. For example, many meteorological models contain parametrizations of sub-grid-scale processes (known as the *physics* in the meteorological literature), which include such things as clouds, precipitation and surface drag. The schemes used to represent these processes can be highly complex and often in-

[2] The term *automatic differentiation* refers to the approach itself, not just to the automatic tools.

clude non-differentiable functions, such as on-off switches. While it is possible for automatic differentiation to deal with such functions it is usually felt that this level of compexity is not necessary in the TLM and adjoint model. Hence a series of simpler parametrizations have been developed solely for use in incremental 4D-Var, that capture the main behaviour of the more complex schemes [100], [51], [82], [74].

An alternative approach, devised by the Met Office, is to start from the premise that the linear model must evolve finite and not infinitesimal perturbations and so there is no need for the linear model to be tangent to any nonlinear model. In this approach the linear model is designed with this in mind. In particular, the resolved dynamics is approximated by a discretization of the linearized continuous equations, with various simplifications in the equations and the discretization. Then simplified parametrizations can be used to represent sub-grid-scale processes [72], [60]. The adjoint model is derived from this approximate linear model by the process of automatic differentiation, ensuring that it provides the exact gradient of the discrete linear cost function.

An essential part of the development of the linear and adjoint models is their testing, as any small mistakes could lead to lack of convergence of the minimization algorithms. Robust tests exist to check the coding of a TLM and adjoint model. The test for the TLM is based on comparing the evolution of a perturbation in the TLM with the evolution of the same perturbation in the nonlinear model. A Taylor series expansion of the nonlinear model operator shows that the evolutions should be closer together as the perturbation size is reduced [81], [65]. Where an inexact TLM is used then this test is not able to differentiate between small coding errors and the desired inexactness. In this case other more subjective tests must be performed [60]. The adjoint model code can be tested by a verification of the adjoint identity

$$(\mathbf{M}_i \delta\mathbf{x}_i)(\mathbf{M}_i \delta\mathbf{x}_i)^T = \delta\mathbf{x}_i^T(\mathbf{M}_i^T(\mathbf{M}_i \delta\mathbf{x}_i)), \tag{9}$$

which should hold to machine precision [77]. For large codes each of these tests should be available for each subroutine as well as at higher levels. A further test, also based on a Taylor expansion, is used to verify that the gradient of the cost function has been correctly coded [77].

3.2 Background error covariances

The background error covariance matrix $\mathbf{B}$ plays a very important role in variational data assimilation. The calculation of these covariances for the assimilation system is a hugely complex task and very dependent on the specific system being modelled. Here we are only able to give an outline of the main steps involved. For further details in the context of atmospheric data assimilation the reader is referred to the comprehensive two-part review article of Bannister [4], [5].

The analysis increment, defined to be the difference between the analysis and background fields, can be shown to lie in the subspace spanned by the columns of the matrix $\mathbf{B}$ [4]. As a consequence the relationships between different entries of

this matrix determine the relationships between increments to different physical variables or between increments at different spatial points. Thus this matrix is fundamental in allowing information to be inferred about unobserved variables or unobserved regions. However, it is usually impossible to represent this matrix in matrix form. If the state vector is of size n then the matrix $\mathbf{B}$ is of size $n \times n$ and when n is of order 10^8 this matrix is impossible to calculate or store. Instead the action of this matrix is usually represented by a variable transform.

We consider the variable transform in the context of incremental variational data assimilation, since that is how it is usually implemented. We define a new variable $\delta\mathbf{z}_i \in \mathbb{R}^n$ and a transformation matrix $\mathbf{U}_i \in \mathbb{R}^{n\times n}$, such that

$$\delta\mathbf{x}_i = \mathbf{U}_i\delta\mathbf{z}_i, \qquad i = 0, \ldots, N. \tag{10}$$

In terms of this new variable the incremental cost function (7) can be written

$$\begin{aligned}\tilde{\mathcal{J}}^{(k)}(\delta\mathbf{z}_0^{(k)}) &= \frac{1}{2}(\delta\mathbf{z}_0^{(k)} - [\mathbf{z}^b - \mathbf{z}_0{}^{(k)}])^{\mathrm{T}}\mathbf{U}_0^T\mathbf{B}^{-1}\mathbf{U}_0(\delta\mathbf{z}_0^{(k)} - [\mathbf{z}^b - \mathbf{z}_0{}^{(k)}]) \\ &+ \frac{1}{2}\sum_{i=0}^{N}(\mathbf{H}_i\mathbf{U}_i\delta\mathbf{z}_i^{(k)} - \mathbf{d}_i^{(k)})^{\mathrm{T}}\mathbf{R}_i^{-1}(\mathbf{H}_i\mathbf{U}_i\delta\mathbf{z}_i^{(k)} - \mathbf{d}_i^{(k)}). \end{aligned} \tag{11}$$

If the variables $\delta\mathbf{z}$ are chosen in such a way that they are uncorrelated then they have identity covariance matrix by definition and so $\mathbf{U}_0^T\mathbf{B}^{-1}\mathbf{U}_0$ can be replaced with the identity in the cost function (11). In this case the cost function no longer contains the original background error covariance matrix; instead it is implicitly defined through the variable transform, with $\mathbf{B} = \mathbf{U}_0\mathbf{U}_0^T$. Furthermore, this variable transform is expected to lead to a better conditioned problem. To understand this we note that the Hessian of the transformed problem is given by

$$\mathbf{I} + \sum_{i=0}^{N}\mathbf{U}_i^T\mathbf{M}(t_i, t_0)^T\mathbf{H}_i^T\mathbf{R}_i^{-1}\mathbf{H}_i\mathbf{M}(t_i, t_0)\mathbf{U}_i, \tag{12}$$

where

$$\mathbf{M}(t_i, t_0) = \mathbf{M}_{i-1}\mathbf{M}_{i-2}\ldots\mathbf{M}_0 \tag{13}$$

is the tangent linear model solution operator from time t_0 to time t_i. Usually the number of observations is less than the number of state variables being estimated and so the Hessian (12) is equal to the identity plus a low rank matrix. Hence it has a minimum eigenvalue equal to one and the condition number (in the two-norm) is equal to the largest eigenvalue.

Of course this theory all relies on being able to choose appropriate variables $\delta\mathbf{z}$ that are truly uncorrelated and it is here that a knowledge of the physical problem is necessary. In presenting how the transform is designed in practice it is easier to think about it in terms of the inverse transform, from model variables to uncorrelated variables. A common approach in numerical weather prediction is to split the inverse transform into two parts. The first part, known as the parameter transform, transforms to physical variables that are assumed

to be uncorrelated between themselves, but still contain spatial correlations. The spatial transform then removes spatial correlations between the physical variables. We now consider each of these transforms in turn.

3.2.1 Parameter transform

In designing a suitable transform of parameters it is necessary to have an understanding of the particular system being modelled, in order to decide which variables are likely to be uncorrelated. For atmospheric models the transform is based on the concept of balanced variables. Balance relationships are diagnostic relationships that exist between certain atmospheric variables. For example, in mid-latitudes and at large horizontal length scales the horizontal wind is approximately in balance with the gradient of the pressure field, through the relationship of geostrophic balance. This relationship can be used in the parameter transform by assuming that the balanced part of the flow is uncorrelated with the unbalanced part [5]. This can be justified by an eigenanalysis of the linearized equation set, which shows that the balanced flow can be associated with one eigenvector and the unbalanced flow with the remaining eigenvectors. Hence under linear evolution these will evolve independently.

The variable that best represents the balanced flow in the atmosphere is potential vorticity (PV) [46] and so it would be natural to use this variable as the basis for the parameter transform. However the transform from PV to the original model variables requires the solution of a three-dimensional elliptic equation as part of the application of the operators $\mathbf{U}_i$. For certain dynamical regimes the PV is well-approximated by the vorticity, which only requires the solution of a two-dimensional equation. Hence early work in this area proposed a transform based on this variable [80] and this is still the basis of the parameter transform in many operational weather forecasting systems [5]. It is recognized that this approximation is not valid in all parts of the atmosphere and it has been demonstrated on simple systems that significant correlations can remain between the transformed variables [53]. For this reason attempts are being made to implement a transformation based on PV in large-scale systems [20], [6].

A similar approach may be followed in other applications, for example in ocean forecasting, though here there has been less work on the design of appropriate transforms than in the meteorological context. In many cases it may be assumed that model variables such as salinity and temperature are uncorrelated and only the spatial transform is needed [99], but work on defining balance relationships has allowed multi-variate covariances to be introduced [97].

3.2.2 Spatial transform

Once the parameter transform has been performed it is assumed that the resulting variables are uncorrelated between themselves. At this point it is necessary to specify the autocovariance information for each parameter through the spatial transform. In atmospheric models it is common to use a Fourier transform in the horizontal and assume that the errors in each wavenumber are uncor-

related. For the vertical correlations a transformation to the eigenvectors of a vertical error covariance matrix is used, with the assumption that the errors associated with each eigenvector are uncorrelated. A scaling transformation is also needed to ensure that the variance of the transformed variables is equal to one. It is often assumed that the correlations are homogeneous (independent of horizontal position) and isotropic (independent of orientation), which simplifies the transformation matrices that must be defined.

In ocean models the complex boundaries near the coast prohibit the simple use of a Fourier transform in the horizontal and so other methods must be used to represent spatial correlations. For example, the application of a correlation operator can be shown to be equivalent to the integration of an appropriately-constructed diffusion equation [96]. This can be used to design correlation models for use in data assimilation systems with irregular boundary conditions [99], [98].

The use of transforms for spatial covariances requires the specification of correlation lengthscales and variances for each of the transformed variables. Since the background field is usually a short-term forecast, these statistics must represent the structure of errors in the forecasting system being used and so be diagnosed from that. An early method for obtaining these statistical parameters used the difference between the observations and the background field (known as the *innovations*) [45]. However, a disadvantage of this method is that it relies on having a sufficient number of observations and is therefore biased towards data-dense areas. The most popular method in atmospheric data assimilation is that known as the 'NMC method' [80][3]. In this method the difference between two different forecasts valid at the same time (for example, a 24-hour and a 48-hour forecast) is taken as a proxy for forecast errors and statistics are taken over a sample of many such forcasts. This method is far from perfect and usually some post-processing of the statistics is necessary to make them realistic (e.g. [49]). A key assumption in this method is that the error covariance matrix represents a statistical average over time. The computational expense of calculating these statistics means that the matrix is kept constant from day to day, perhaps with different statistics being used with a change of season. More recently there has been interest in developing methods for estimating these statistics that vary from day to day, since it is expected that the actual background errors will depend on the underlying flow. Such flow-dependent statistics arise naturally in ensemble methods of data assimilation, such as the ensemble Kalman filter. Methods are currently being designed to obtain some flow-dependent information in variational assimilation, by combining information from ensembles of forecasts with the statistically-averaged error covariance matrix, for example [13].

[3]So called because it was first introduced in the National Meteorological Center of the USA, now the National Center for Environmental Prediction.

3.3 Observation errors

As well as representing the errors in the background field it is important to treat properly the errors in the observations within a variational data assimilation system. Observational data received into operational weather and ocean forecasting centres can contain errors from a variety of sources, including limitations in the measuring instrument, biases in the measurements and errors simply due to human error in recording the measurement. The theory of variational data assimilation assumes that all observational errors are random, unbiased errors with a Gaussian distribution and known covariance. It is therefore important that as many of these sources of error as possible are accounted for in the data assimilation system.

A first essential step in an operational data assimilation system is to perform a quality control check on the data themselves. This may consist of several stages. First a check for obvious errors is made, so that if, for example, a ship observation is reported over a land point it will be rejected from the assimilation. Then a so-called 'background check' may be made to see how close the observation is to the forecast background field. If the difference from the background is too large when compared with its expected error variance then the observation may be rejected and not used in the assimilation [1]. Once this check has been performed the next step is to identify observations that may have gross errors. This can be done either outside or within the assimilation process. Outside the assimilation each observation can be checked against nearby observations and any observations that largely disagree with others can be rejected [83]. Alternatively this check can be included in the assimilation process using the variational quality control method [50], [1]. In this method the observation errors are assumed to be a weighted combination of a standard Gaussian distribution and a flat probability distribution function, with the weights determined by the probability of gross error of the observation. In this way observations that have a high probability of gross error are given very little weight in the analysis. These probabilities are updated during the iteration procedure by comparison with the current estimate of the state, to allow observations to be given more or less weight as the assimilation progresses.

A second important aspect of observation errors is the treatment of systematic errors, or biases, in the observations. This is particularly important for satellite radiance data, where biases may occur from changes in the measuring instrument over time or from errors in the radiative transfer model needed as part of the observation operator [42]. Since the assimilation scheme assumes that the observations are unbiased, any biases in the observations can introduce biases into the analyses. As with the quality control these biases may be treated offline or within the assimilation scheme. For each satellite channel a bias model is assumed in such a way that we can define a new observation operator for the biased measurement

$$\tilde{\mathcal{H}}(\mathbf{x}, \boldsymbol{\beta}) = \mathcal{H}(\mathbf{x}) + \mathbf{b}(\boldsymbol{\beta}, \mathbf{x}), \tag{14}$$

with

$$\mathbf{b}(\boldsymbol{\beta}, \mathbf{x}) = \sum_{j=0}^{N_p} \beta_j \mathbf{p}_j(\mathbf{x}), \tag{15}$$

where $\mathbf{p}_j$ are predictors for $j = 0, \ldots, N_p$ and β_j are scalar coefficients [25]. A few predictor states are chosen that may be related to the state at the observation positions. The coefficients $\boldsymbol{\beta}$ can then be estimated in an offline regression using a few weeks of data [42] or a variational procedure can be used to estimate these coefficients. This can also be included directly in the assimilation procedure by including (14) in the cost function in place of the standard observation operator and including a background estimate $\boldsymbol{\beta}^b$ of $\boldsymbol{\beta}$ with covariance $\mathbf{B}_{\boldsymbol{\beta}}$. The 4D-Var assimilation problem is then to minimize

$$\begin{aligned} \mathcal{J}_{\boldsymbol{\beta}}(\mathbf{x}_0, \boldsymbol{\beta}) &= \frac{1}{2}(\mathbf{x}_0 - \mathbf{x}^b)^{\mathrm{T}} \mathbf{B}^{-1} (\mathbf{x}_0 - \mathbf{x}^b) + \frac{1}{2}(\boldsymbol{\beta} - \boldsymbol{\beta}^b)^T \mathbf{B}_{\boldsymbol{\beta}}^{-1} (\boldsymbol{\beta} - \boldsymbol{\beta}^b) \\ + \quad & \frac{1}{2} \sum_{i=0}^{N} (\mathcal{H}_i(\mathbf{x}_i) + \mathbf{b}(\boldsymbol{\beta}, \mathbf{x}_i) - \mathbf{y}_i)^{\mathrm{T}} \mathbf{R}_i^{-1} (\mathcal{H}_i(\mathbf{x}_i) + \mathbf{b}(\boldsymbol{\beta}, \mathbf{x}_i) - \mathbf{y}_i) \end{aligned} \tag{16}$$

subject to the dynamical equations, to estimate the state $\mathbf{x}_0$ and the coefficients $\boldsymbol{\beta}$ simulteneously [25]. Alternatively, a variational procedure can be used to estimate these coefficients offline at regular intervals, using the previous value as the background for the new estimate [2].

Finally we consider the specification of the observation error covariance matrix, which represents the covariance of the random components of the observation error. It is important to note that this error is defined by the difference between the actual measurement and the model representation of the true state $\mathbf{x}^t$, i.e. the error $\boldsymbol{\epsilon}_i^o$ at time t_i is given by

$$\boldsymbol{\epsilon}_i^o = \mathbf{y}_i - \mathcal{H}_i(\mathbf{x}_i^t). \tag{17}$$

This means that the error includes different components arising from the accuracy of the measuring instrument (instrument error), errors in the observation operator $\mathcal{H}_i$ and errors due to the difference in spatial resolution between the measurement and the model state (known as representativity error). The instrument error is the easiest to treat, since the variances of this error can usually be obtained from the instrument manufacturer and it is normally safe to assume that these errors are uncorrelated. However, this may not always be the case. For example, measurements derived by preprocessing satellite data may include spatial correlations [12]. Errors in the observation operator may include such things as errors in the radiative transfer models used to model satellite data, which can lead to error correlations between different satellite channels [11], [88].

Although it is recognized that observation error correlations exist, particularly with respect to satellite data, the correlations are not usually very well treated in current operational forecasting systems. Often the correlations are ignored and it is assumed that the observation error covariance matrix is diagonal. To balance this assumption either the error variances are inflated [44] or the

data are thinned so that many fewer of them are used [24]. The reasons for this are the difficulty in calculating what the error correlations should be and the difficulty in then representing these correlations within an assimilation scheme in a way that the inverse correlation matrix can easily be applied. To estimate the correlations in satellite data the methods that have mainly been used are a comparison with independent measurements from radiosondes, based on the method of [45], and the use of diagnostics calculated from the data assimilation system itself, based on [26]. Various ways of then representing these correlations within the data assimilation system have been proposed, including the use of a circulant matrix [43], an eigenvalue decomposition [27] and a Markov matrix [88]. However there is so far little use of these methods in operational practice.

3.4 Optimization methods

The minimization of the inner loop cost function (7) requires the use of a suitable optimization algorithm. For the large problems of environmental modelling there are two particularly important constraints. The first is that because of the number of variables in the system it is not possible to obtain second derivative information. The Hessian or second derivative matrix would contain of the order 10^{16} elements, which is impossible to calculate or to store. Hence only methods that require first derivative information can be used. The second constraint is that often these problems must be solved within a real-time forecasting system and hence the computer time that can be used to solve the problem is very limited. Hence the methods much use as few function evaluations as possible. This means that usually the problem is not allowed to run to full convergence and the use of any line search algorithms is prohibitively expensive. Traditionally the algorithms that have most been used within data assimilation systems are quasi-Newton algorithms and conjugate gradient or related Lanczos algorithms. The mathematical details of these algorithms are well explained elsewhere (e.g. [78]) and so here we limit discussion to their implementation in data assimilation systems.

An essential aspect of the minimization procedure for variational data assimilation is an appropriate preconditioning. Experimental evidence indicates that that the Hessian of the inner loop cost function (7) is badly conditioned and that this arises from the ill-conditioning of the background error covariance matrix [69]. This has been further confirmed by theoretical results that bound the condition number of the Hessian of the cost function in terms of the condition number of this covariance matrix [39], [40]. The first level of preconditioning that is applied is therefore to transform the problem to new variables, as described in section 3.2. The transformed problem (11) can be shown in general to be better conditioned both in theory and in practice [69], [31], [39], [40]. However, even after this transformation the problem is not very well-conditioned and can have a condition number of order $10^3 - 10^4$ [29], [41]. Experiments in the ECMWF system showed that the ill-conditioning that remains is related to the inclusion of dense, accurate surface observations over Europe [93] and this has also been shown to be true for the system of the Met

Office [41]. This can be explained by theoretical bounds obtained by [39], [41] that show that the condition number of the transformed problem increases as the spacing between observations decreases and as observations become more accurate. Hence ideally a second level of preconditioning is required after the variable transformation has been performed.

In order to implement a further preconditioning some knowledge of the Hessian (12) of the transformed cost function is required. One way that this can be obtained is by using a Lanczos algorithm to perform the inner loop minimization. The Lanczos method produces estimates of the leading eigenvectors and eigenvalues of the Hessian of the function being minimized. If the first m eigenvalues λ_j and eigenvectors $\mathbf{u}_j, j = 1, \ldots, m$ have sufficiently converged then the Hessian (12) can be approximated by the expression

$$\mathbf{I} + \sum_{j=1}^{m} (\lambda_j - 1)\mathbf{u}_j \mathbf{u}_j^T. \tag{18}$$

This expression can then be used for the preconditioning of subsequent minimizations, under the assumption that the Hessian does not change greatly between one minimization and another [29], [94]. This method, known as spectral preconditioning, is used in the operational forecast system of ECMWF, where three outer loops are performed for each assimilation. During the first inner loop minimization the Lanczos vectors are stored and these are then used to precondition the mininimization of the second and third inner loop cost functions [29]. It has been shown that this preconditioner belongs to a larger class of limited memory preconditioners [94]. The authors of [94] propose an alternative preconditioner from the same class, based on the Ritz pairs of the Hessian. They found that this can provide an improvement over spectral preconditioning when the estimates of the Hessian eigenpairs are inaccurate. A similar result was also found in the Regional Ocean Modelling System (ROMS), in which both of these preconditioners are implemented [76]. One drawback of both of these methods is that, in order to generate the required information, the first minimization must be performed before any preconditioning can be applied. So far little attention has been paid to preconditioning of this first minimization.

With any minimization method it is important to specify appropriate stopping criteria and this is also the case in variational data assimilation. It has been proved that the inner-loop steps of the Gauss-Newton method need to be solved to sufficient accuracy in order to ensure convergence of the outer loops [34]. The theory has been used to show how it is natural to use an inner-loop stopping criterion based on the relative change in the norm of the gradient [59]. The tolerance used to stop the iterations must therefore be chosen carefully. If it is too high then there is no guarantee that the outer loop steps will converge. However the convergence should not be pushed below the level of noise on the observations, as then small spatial scales are adjusted to fit the observational noise [55]. In many practical forecasting problems such care is not always taken and other criteria are introduced. There are two main reasons for this. One is that in a time-critical forecasting system it may considered more important

to solve each minimization problem using approximately the same amount of wall-clock time rather than to the same accuracy. The second reason is that the preconditioning techniques decribed in this section require a minimum number of iterations to be performed on the first inner-loop minimization in order to acquire sufficiently accurate information about the Hessian. Hence criteria that have been introduced include stopping the iterations when the value of the cost function is close to its expected minimum value [70] or using a fixed number of iterations, particularly for the first minimization [93].

3.5 Reduced order approaches

As was mentioned in section 2.1 a major advantage of the incremental approach is that the inner loop problem may be solved in a smaller dimensional space than the outer loop update of the linearization trajectory. Within environmental prediction lower spatial resolution systems have often been used in the inner loop step, with the full resolution nonlinear model being used in the outer loop. Further simplifications may also be made to the linear dynamical model used in the inner loop, such as using simplified parametrizations of sub-grid scale processes as described in section 3.1. While a change in resolution is certainly the simplest way to achieve a more computationally tractable inner loop problem, it does not necessarily provide the most accurate low order representation of the system. In order to improve on this other reduced order approaches have been investigated in the context of incremental 4D-Var. These essentially fall into two categories, methods based on principal component analysis and methods based on near-optimal reduction of dynamical systems.

Principal component analysis, which is often referred to as principal orthogonal decomposition (POD) or the method of empirical orthogonal functions (EOFs), aims to represent the solution of the assimilation problem as a linear combination of basis vectors. The basis vectors are chosen to represent the leading directions of variability in the model and are calculated using a series of model states, or 'snapshots', from an integration of the nonlinear model. Such a method was used in an ocean model assimilation by [84]. From the sample of model states the authors generate the matrix $\mathbf{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_l)$, where $\mathbf{X}_i$ is the difference between the model state at time t_i and the mean state. The covariance matrix $\mathbf{X}^T\mathbf{X}$ is then diagonalised to find a set of orthonormal eigenvectors $\mathbf{v}_i$ (EOFs) and associated eigenvalues $\lambda_i, i = 1, \ldots, l$. The solution $\delta\mathbf{x}_0$ to the inner loop minimization problem (7) is then defined by an expansion of the leading r eigenvectors

$$\delta\mathbf{x}_0 = \sum_{i=0}^{r} w_i \mathbf{v}_i = \mathbf{V}\mathbf{w} \tag{19}$$

where $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_r)$ is the matrix of the leading r eigenvectors and $\mathbf{w} = (w_1, \ldots, w_r)^T$ are the weights to be determined. In this case the matrix $\mathbf{V}$ acts as a variable transformation in a similar way to the parameter transform (10)

and so the background term can be written in the form

$$\mathcal{J}_b(\mathbf{w}) = \frac{1}{2}\mathbf{w}^T\mathbf{B}_w^{-1}\mathbf{w}, \tag{20}$$

where the covariance matrix $\mathbf{B}_w$ is taken to be the diagonal matrix of eigenvalues. This method has been applied to assimilation in ocean models in an idealized setting [84] and using real data [47]. It is noted that the assumption behind this method is that the variability of the system can be well described by a low dimensional space. Although the approach reduces the size of the space in which the minimization is performed, the tangent linear model (8) must still be integrated at full resolution on each iteration.

An alternative approach, based on POD, was put forward by [17]. In that work the solution to the full nonlinear 4D-Var problem is expressed as a perturbation from the sample mean that is expanded in terms of basis functions Φ_i. These basis functions are derived in a similar way to the EOFs, but by then projecting the perturbation fields $\mathbf{X}$ onto the eigenvectors $\mathbf{v}_i$. In this work the authors solve the nonlinear 4D-Var cost function (3) in the reduced space. As well as expressing the background term in terms of coefficients of the basis functions they also derive a Galerkin projection of the dynamical model onto the basis functions for use in the observation term. Thus this formulation has the advantage that the dynamical model and its adjoint are also expressed in the reduced space. Again this method relies on the snapshots being able to capture a low-dimensional subspace that adequately describes the full system.

A disadvantage with both the EOF and POD methods is that they do not use any information about the data assimilation problem itself within the reduction procedure. There have been two approaches proposed to improve on this. The first is an adaption of the POD method, called dual-weighted POD. In this method the snapshot perturbations $\mathbf{X}$ are weighted according the sensitivity of the cost function at the time of the snapshot, where the weights are calculated using the adjoint model [22]. The other approach, put forward in the series of papers [62], [61], [9], is to use near-optimal model order reduction methods for linear dynamical systems to derive a reduced order model and observation operator. The inner loop problem of incremental 4D-Var (7) is subject to the dynamical system described by the evolution equation (8) and the output equation

$$\mathbf{d}_i = \mathbf{H}_i\delta\mathbf{x}_i. \tag{21}$$

Model reduction seeks linear restriction operators $\mathbf{S}_i^T$ and prolongation operators $\mathbf{T}_i$ that map the perturbation $\delta\mathbf{x}_i \in \mathbb{R}^n$ to $\delta\hat{\mathbf{x}}_i \in \mathbb{R}^r$ with $r << n$. These operators are chosen such that the output of the projected system

$$\delta\hat{\mathbf{x}}_{i+1} = \mathbf{S}_i^T\mathbf{M}_i\mathbf{T}_i\delta\hat{\mathbf{x}}_i \tag{22}$$

$$\hat{\mathbf{d}}_i = \mathbf{H}_i\mathbf{T}_i\delta\hat{\mathbf{x}}_i \tag{23}$$

approximates well the full dynamical system. The inner loop problem can then

be defined in the reduced space as the minimization of

$$\begin{aligned}\min \hat{\mathcal{J}}^{(k)}[\delta\hat{\mathbf{x}}_0^{(k)}] &= \frac{1}{2}(\delta\hat{\mathbf{x}}_0^{(k)} - \mathbf{S}^T[\mathbf{x}^b - \mathbf{x}_0{}^{(k)}])^{\mathrm{T}} \\ &\times (\mathbf{S}_i^T\mathbf{B}\mathbf{S}_i)^{-1}(\delta\hat{\mathbf{x}}_0^{(k)} - \mathbf{S}_i^T[\mathbf{x}^b - \mathbf{x}_0{}^{(k)}]) \\ &+ \frac{1}{2}\sum_{i=0}^{N}(\mathbf{H}_i\mathbf{T}_i\delta\hat{\mathbf{x}}_i^{(k)} - \mathbf{d}_i^{(k)})^{\mathrm{T}}\mathbf{R}^{-1}(\mathbf{H}_i\mathbf{T}_i\delta\hat{\mathbf{x}}_i^{(k)} - \mathbf{d}_i^{(k)}),\end{aligned}$$

subject to the reduced dynamical model (22). The linearization state is then updated with the perturbation

$$\delta\mathbf{x}_0^{(k)} = \mathbf{T}_i\delta\hat{\mathbf{x}}_0^{(k)}. \tag{24}$$

The authors of these papers use the method of balanced truncation to demonstrate this method in the case where the operators $\mathbf{M}$ and $\mathbf{H}$ are time-invariant. In idealised models they show how this method improves the solution with respect to using low resolution models and how it is important to use information about the assimilation problem in the reduction procedure, including information about the background and observation error covariance matrices. However, whereas reduction methods based on POD can be implemented in large systems, the method of balanced truncation cannot. Efforts are being made to design near-optimal reduction methods for large systems based on Krylov methods [16], but these methods have not yet been tried out in data assimilation for large systems.

3.6 Issues for nested models

For very high resolution weather and ocean forecasting operational centres often use models covering only the domain of interest that are nested in a larger model, often of lower resolution, which we refer to here as the parent model. In most of the systems the nesting is a one-way nesting, whereby lateral boundary conditions for the nested model are provided by the parent model, but there is no feedback from the high resolution nested model to the parent model. This presents particular challenges for the application of variational data assimilation. For problems specific to high resolution weather forecasting we refer the reader to the review articles [79] and [23]. Here we consider only more general problems arising from using a high resolution nested grid, in particular treatment of the lateral boundary conditions and of the difference in representation of spatial scales between the parent and nested models.

With respect to the lateral boundary conditions, a decision must be made as to whether to estimate them as part of the assimilation procedure or to assume that they do not change. Both approaches have been used in practice. In the operational weather forecasting system of the Met Office the lateral boundary conditions are not updated, but are fixed by the parent model. Hence the increment $\delta\mathbf{x}$ on the boundary is set to zero. This has advantages for the practical implementation of the scheme. In particular it allows a simple sine

transform to be used in the definition of the spatial background error covariances described in section 3.2, which then enforces zero boundary increments [69]. However, observational information close to the boundaries can be difficult to use, since the nested model cannot use observations lying outside the domain and the analysis inside the domain may not be consistent with the boundary conditions provided [3], [36]. This can lead to features being artificially cut-off close to the boundaries.

The alternative approach is to estimate the boundary variables within the assimilation procedure [37], [54], [38]. In this way observations inside the nested domain can update the boundary values and so it is possible to ensure that the analysis is consistent throughout the domain. However in this case it is no longer possible to apply a sine transform to impose the spatial background error covariances. In order to be able to apply a spectral transformation an extension zone is created around the domain to obtain fields that are horizontally periodic. A Fourier transform can then be applied. One difficulty in analysing the boundaries in this way is that the lateral boundary conditions are only updated during the assimilation period. During the subsequent forecast no updates are available and the values from the parent model must be used, so there is some inconsistency between the boundary conditions of the analysis and those of the forecast.

The second challenge we consider is the difference in the spatial scales that can be represented in the nested and parent models. In particular, since the nested model often covers only a small domain, the assimilation scheme is not able to analyse adequately scales of the size of the domain and larger. In applications such as weather prediction it is important to capture these larger scales, since the physical system is inherently multiscale, with strong feedbacks between large and small scales. Hence attempts have been made to improve the large-scale information in nested model data assimilation by providing information on these scales from a parent model analysis. For example, the Met Office experimented with a system that combined large scale increments from a parent model analysis with the small scale increments from the nested model analysis [3]. In this method the large scales of the nested model analysis are forced to be equal to those of the parent model. An alternative, proposed by [36], is to use the large scales of the parent analysis over the nested model domain as a weak constraint on the variational problem. This is done by adding an extra term to the inner loop cost function (7) that measures the distance between the large scales of the global analysis and those forecast by the nested model. This means that the analysis is constrained by large scales from the parent model, through this additional term, and by large scales from the nested model, through the background term. In theory this should introduce another term including the cross-correlation between these two sources of information. However, in their demonstration of the method in a 3D-Var scheme of the ALADIN model at Météo-France the authors of [36] concluded that this correlation could be neglected, though at the cost of some inaccuracy.

A more theoretical study of this problem was carried out by [7]. They used a spectral analysis to show how information from waves longer than the domain

size is projected onto different scales in the nested model domain, corresponding to the lowest wave numbers that can be represented on this domain. They demonstrated that by giving more weight to these scales in the background term of the cost function it was possible to retain more of the large scale information from a parent model background. In this method only the large spatial scales from the parent model are used as a constraint in the assimilation, as in [3], but they are not imposed exactly and may be altered by the assimilation process. The authors of [7] demonstrated benefit from this in an idealised system, but the method has not been tested in a realistic model.

3.7 Weak constraint variational assimilation

The formulation of variational data assimilation presented in section 2 assumes that the discrete dynamical model (1) is an exact representation of the physical system being observed. In practice we know that the models contain errors, caused by limitations in our knowledge of the physical equations and limitations in the numerical modelling, such as the need for sub-grid scale parametrizations. In theory it is possible to account for and estimate such errors in variational data assimilation, though implementation in practice is more complicated. We assume an additive error to the model equations, so that the true dynamical system can be written

$$\mathbf{x}_{i+1} = \mathcal{M}_i(\mathbf{x}_i) + \boldsymbol{\eta}_i, \tag{25}$$

where $\boldsymbol{\eta}_i$ is the unknown model error at time t_i. Then we can define a weak constraint 4D-Var problem, in which the model equations do not have to be exactly satisfied over the assimilation window. We define a cost function of the form

$$\begin{aligned} \mathcal{J}(\mathbf{x}_0, \boldsymbol{\eta}_0, \ldots, \boldsymbol{\eta}_{N-1}) &= \frac{1}{2}(\mathbf{x}_0 - \mathbf{x}^b)^{\mathrm{T}}\mathbf{B}^{-1}(\mathbf{x}_0 - \mathbf{x}^b) \\ + \quad \frac{1}{2}\sum_{i=0}^{N}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i)^{\mathrm{T}}\mathbf{R}_i^{-1}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i) &+ \frac{1}{2}\sum_{i=0}^{N-1}\boldsymbol{\eta}_i^T\mathbf{Q}_i^{-1}\boldsymbol{\eta}_i \end{aligned} \tag{26}$$

subject to (25), where $\mathbf{Q}_i$ is the covariance matrix associated with the model errors $\boldsymbol{\eta}_i$. The weak constraint problem is then to minimize (26) with respect to the initial state $\mathbf{x}_0$ and all the model errors $\boldsymbol{\eta}_i$.

An alternative formulation of the weak constraint problem (26) is to write it in terms of the model state $\mathbf{x}_i$ at each time t_i rather than in terms of the model errors. This leads to the cost function

$$\begin{aligned} \mathcal{J}(\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_N) &= \frac{1}{2}(\mathbf{x}_0 - \mathbf{x}^b)^{\mathrm{T}}\mathbf{B}^{-1}(\mathbf{x}_0 - \mathbf{x}^b) \\ &+ \frac{1}{2}\sum_{i=0}^{N}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i)^{\mathrm{T}}\mathbf{R}_i^{-1}(\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i) \\ &+ \frac{1}{2}\sum_{i=0}^{N-1}(\mathbf{x}_{i+1} - \mathcal{M}_i(\mathbf{x}_i))^T\mathbf{Q}_i^{-1}(\mathbf{x}_{i+1} - \mathcal{M}_i(\mathbf{x}_i)), \end{aligned} \tag{27}$$

which is minimized subject to (25). Although (27) is mathematically equivalent to (26), in practice their implementations may be different. In [92] both formulations were presented in the incremental version of 4D-Var as possibilities for inclusion in the ECMWF system.

The inclusion of the model errors at each observation time increases the size of the argument of $\mathcal{J}$ by a factor of $N+1$, the number of observation times. One way to reduce this cost is by assuming a relationship in time between the model errors $\boldsymbol{\eta}_i$. Theoretical work by [35] used an augmented state approach to solve for the state and the model error, with a dynamical equation used to explain the evolution of the error. They introduced a general form for the error evolution, including both a systematic and random component of the error. Various options for the systematic evolution were proposed, including a contant bias error and simple dynamical evolutions, and the methods were illustrated on simple systems. In the context of a regional atmospheric model [101] demonstrated a weak-constraint 4D-Var system under the assumption that the model error was serially correlated and obeyed a first order Markov process.

Since this early work there have been several idealised studies with weak-constraint 4D-Var, but the move towards operational implementations in large-scale systems has been slow. One of the biggest challenges remains the specification of the model error covariance matrix $\mathbf{Q}_i$ for real systems. An initial idea was to take this matrix to be a scalar multiple of the background error covariance matrix $\mathbf{B}$. However, in experiments with the ECMWF atmospheric forecasting system using formulation (26), [93] showed that this choice implies that corrections to the model error lie in the same space as those to the background. This leads to estimates of model error that are very similar to the increments to the initial conditions. An alternative method, proposed in the same paper, is to estimate the statistics of $\mathbf{Q}_i$ from an ensemble of differences between model tendency fields using the NMC method, in a similar way that differences between the model fields themselves are used in the the estimation of the background error covariances (as explained in section 3.2). [93] interprets differences between these tendencies as a proxy for the uncertainty in the model forcing. The statistics from this sample are then fit to the same statistical model as is used for the matrix $\mathbf{B}$. The use of a covariance matrix estimated in this way was shown to give an improvement in weak-constraint 4D-Var experiments that assumed a constant error over the assimilation window.

The work of [66] illustrated the implementation of weak-constraint 4D-Var using such a matrix, again in the ECMWF system, to estimate a constant bias error in the stratosphere, where the model is known to have biases. A similar scheme has been introduced into the operational assimilation system of ECMWF [30]. In this implementation the deviation of the error from its mean value is minimized, so that the last term of (26) becomes

$$\frac{1}{2}(\boldsymbol{\eta}-\bar{\boldsymbol{\eta}})^T\mathbf{Q}^{-1}(\boldsymbol{\eta}-\bar{\boldsymbol{\eta}}), \tag{28}$$

where $\bar{\boldsymbol{\eta}}$ is the estimate of the model bias from the previous analysis cycle. In this way the assimilation ensures that the estimated error does not vary too

quickly from one analysis cycle to the next.

Despite these initial successes much more work is needed. One particular difficulty is that it is not clear how to differentiate between model bias and observation bias, since the assimilation only measures the difference between the model and the observations. [93] showed a case study of observation bias being interpreted as a model error by weak-constraint 4D-Var. This problem was discussed further by [64] in the context of ocean data assimilation. They suggested that to estimate both model and observation bias it is necessary to include information on the spatial and temporal structure of these biases in the covariance matrices.

In order to then move away from the assumption of a constant bias and treat time-varying systematic and random model errors, more sophisticated methods for describing the evolution of errors must be developed. This evolution is likely to be dependent on the specific model being used, yet general methods for representing this are also needed. At the same time efficient and accurate representations of the covariances of these model errors must be found. The use of the weak-constraint formulation of 4D-Var holds much promise to counteract the inadequacies of models, but many challenges remain open to be able to implement this in very large environmental models

4 Summary and future perspectives

Variational data assimilation is now a well-established method for combining observational data with very large environmental models. However, as has been illustrated in this article, its successful implementation requires careful and judicious choices in each aspect of the assimilation scheme. In some cases these choices are determined by the physical system being modelled or the observational data available, such as the specification of the error covariances in the system. In other cases the choices may be determined by the size of the problem and the need to solve it in an efficient manner, often for real-time forecasting, or by features of the numerical model itself, such as lateral boundary conditions. In each instance the choices to be made will inevitably be a compromise between the ideal solution and what is practically feasible in a given system. We have presented some of the solutions that have been found that have allowed variational data assimilation to be implemented in large environmental forecasting systems. Nevertheless much research continues to improve on these solutions so as to find better estimates of the state and so produce better forecasts.

One particularly active area in numerical weather prediction is the desire to use more information from ensembles of forecasts to provide time-varying covariances for the background errors, combining the advantages of ensemble methods with the advantages of 4D-Var. ECMWF have implemented a system in which an ensemble of 4D-Var assimilations are run and the statistics from this ensemble are used to update the variances of the background errors [10]. Extensions to this method to calculate also the covariance information are being sought. An alternative approach is to use information from ensembles of

forecasts to calculate covariance information throughout the whole assimilation window. This method was proposed by [67] and tested in a global weather prediction model by [14], [15]. An advantage of this method is that the tangent linear and adjoint models are not required in the 4D-Var, since all the evolution information comes through the ensemble of nonlinear model forecasts. Hence this makes development of the system much easier.

Besides the many great challenges that we have discussed in this article, new challenges are arising for the future evolution of variational data assimilation systems. The advent of massively parallel computers means that the algorithms used currently to solve the assimilation problem may no longer be efficient on future computer architectures. Hence work is needed to develop new algorithms to solve the problem, particularly with respect to efficient minimization and preconditioning methods. This may be easier as systems move to a weak-constraint form of 4D-Var but, as discussed above, that introduces its own difficulties [30]. Another challenge comes from the move towards more integrated Earth-system models, with different environmental models coupled to each other. For example, for seasonal to decadal prediction it is now common to use coupled atmosphere-ocean models, but the initialization of these models with data assimilation is still in its infancy. Particular problems arise from the very different time scales in the atmosphere and ocean system and from the model biases in atmosphere and ocean models. Some work has been done to implement 4D-Var in such systems in order to estimate the ocean state and coupling parameters [89], [75], but the estimation of the complete state in coupled atmosphere-ocean models remains an open problem for the coming years.

References

[1] E. Andersson and H. Järvinen. Variational quality control. *Quart. J. Roy. Meteor. Soc.*, 125(554):697–722, 1999.

[2] T. Auligné, A.P. McNally, and D.P. Dee. Adaptive bias correction for satellite data in a numerical weather prediction system. *Quart. J. Roy. Meteor. Soc.*, 133(624):631–642, 2007.

[3] S. Ballard, Z. Li, M. Dixon, S. Swarbrick, O. Stiller, and H. Lean. Development of 1-4km resolution data assimilation for nowcasting at the met office. In *World Weather Research Program Symposium on Nowcasting and Very Short Range Forecasting (WSN05)*, page Paper 3.02, 2005.

[4] R.N. Bannister. A review of forecast error covariance statistics in atmospheric variational data assimilation. i: Characteristics and measurements of forecast error covariances. *Quart. J. Roy. Meteor. Soc.*, 134(637):1951–1970, 2008.

[5] R.N. Bannister. A review of forecast error covariance statistics in atmospheric variational data assimilation. ii: Modelling the forecast error

covariance statistics. *Quart. J. Roy. Meteor. Soc.*, 134(637):1971–1996, 2008.

[6] R.N. Bannister and M.J.P. Cullen. A regime-dependent balanced control variable based on potential vorticity. In *Proceedings of ECMWF workshop in flow-dependent aspects of data assimilation*, pages 1–13, 2007.

[7] G.M. Baxter, S.L. Dance, A.S. Lawless, and N.K. Nichols. Four-dimensional variational data assimilation for high resolution nested models. *Computers & Fluids*, 46(1):137–141, 2011.

[8] C. Bischof, A. Carle, G. Corliss, A. Griewank, and P. Hovland. ADIFOR: Generating derivative codes from Fortran programs. *Scientific Programming*, 1:11–29, 1992.

[9] C. Boess, A.S. Lawless, N.K. Nichols, and A. Bunse-Gerstner. State estimation using model order reduction for unstable systems. *Computers & Fluids*, 46(1):155–160, 2011.

[10] M. Bonavita, L. Isaksen, and E. Hólm. On the use of eda background error variances in the ecmwf 4d-var. *Quart. J. Roy. Meteor. Soc.*, page doi: 10.1002/qj.1899, 2012.

[11] N. Bormann and P. Bauer. Estimates of spatial and interchannel observation-error characteristics for current sounder radiances for numerical weather prediction. i: Methods and application to ATOVS data. *Quart. J. Roy. Meteor. Soc.*, 136(649):1036–1050, 2010.

[12] N. Bormann, S. Saarinen, G. Kelly, and J.N. Thepaut. The spatial structure of observation errors in atmospheric motion vectors from geostationary satellite data. *Mon. Wea. Rev.*, 131(4):706–718, 2003.

[13] M. Buehner. Ensemble-derived stationary and flow-dependent background-error covariances: Evaluation in a quasi-operational nwp setting. *Quart. J. Roy. Meteor. Soc.*, 131(607):1013–1043, 2005.

[14] M. Buehner, P.L. Houtekamer, C. Charette, H.L. Mitchell, and B. He. Intercomparison of variational data assimilation and the ensemble kalman filter for global deterministic nwp. part i: Description and single-observation experiments. *Mon. Wea. Rev.*, 138(5):1550–1566, 2010.

[15] M. Buehner, P.L. Houtekamer, C. Charette, H.L. Mitchell, and B. He. Intercomparison of variational data assimilation and the ensemble kalman filter for global deterministic nwp. part ii: One-month experiments with real observations. *Mon. Wea. Rev.*, 138(5):1567–1586, 2010.

[16] A. Bunse-Gerstner, D. Kubalinska, G. Vossen, and D. Wilczek. h2-norm optimal model reduction for large scale discrete dynamical mimo systems. *Journal of computational and applied mathematics*, 233(5):1202–1216, 2010.

[17] Y. Cao, J. Zhu, I.M. Navon, and Z. Luo. A reduced-order approach to four-dimensional variational data assimilation using proper orthogonal decomposition. *International Journal for Numerical Methods in Fluids*, 53(10):1571–1583, 2007.

[18] W.C. Chao and L-P. Chang. Development of a four-dimensional variational analysis system using the adjoint method at GLA. Part I: Dynamics. *Mon. Wea. Rev.*, 120:1661–1673, 1992.

[19] P. Courtier, J-N. Thepaut, and A. Hollingsworth. A strategy for operational implementation of 4D-Var, using an incremental approach. *Quart. J. Roy. Meteor. Soc.*, 120:1367–1387, 1994.

[20] M.J.P. Cullen. Four-dimensional variational data assimilation: A new formulation of the background-error covariance matrix based on a potential-vorticity representation. *Quart. J. Roy. Meteor. Soc.*, 129:2777–2796, 2003.

[21] M.J.P. Cullen. A demonstration of 4d-var using a time-distributed background term. *Quart. J. Roy. Meteor. Soc.*, 136(650):1301–1315, 2010.

[22] D.N. Daescu and I.M. Navon. A dual-weighted approach to order reduction in 4dvar data assimilation. *Mon. Wea. Rev.*, 136(3):1026–1041, 2008.

[23] S.L. Dance. Issues in high resolution limited area data assimilation for quantitative precipitation forecasting. *Physica D: Nonlinear Phenomena*, 196(1):1–27, 2004.

[24] M.L. Dando, A.J. Thorpe, and J.R. Eyre. The optimal density of atmospheric sounder observations in the met office nwp system. *Quart. J. Roy. Meteor. Soc.*, 133(629):1933–1943, 2007.

[25] D.P. Dee. Bias and data assimilation. *Quart. J. Roy. Meteor. Soc.*, 131(613):3323–3343, 2005.

[26] G. Desroziers, L. Berre, B. Chapnik, and P. Poli. Diagnosis of observation, background and analysis-error statistics in observation space. *Quart. J. Roy. Meteor. Soc.*, 131(613):3385–3396, 2005.

[27] M. Fisher. Accounting for correlated observation error in the ecmwf analysis. Technical memorandum MF/05106, ECMWF, 2005.

[28] M. Fisher and D.J. Lary. Lagrangian four-dimensional variational data assimilation of chemical species. *Quart. J. Roy. Meteor. Soc.*, 121:1681–1704, 1995.

[29] M. Fisher, J. Nocedal, Y. Trémolet, and S.J. Wright. Data assimilation in weather forecasting: a case study in pde-constrained optimization. *Optimization and Engineering*, 10(3):409–426, 2009.

[30] M. Fisher, Y. Trémolet, H. Auvinen, D. Tan, and P. Poli. Weak-constraint and long-window 4D-Var. Technical memorandum 655, ECMWF, 2011.

[31] P. Gauthier, C. Charette, L. Fillion, P. Koclas, and S. Laroche. Implementation of a 3d variational data assimilation system at the canadian meteorological centre. part i: The global analysis. *Atmosphere-Ocean*, 37(2):103–156, 1999.

[32] P. Gauthier, M. Tanguay, S. Laroche, S. Pellerin, and J. Morneau. Extension of 3dvar to 4dvar: Implementation of 4dvar at the meteorological service of canada. *Mon. Wea. Rev.*, 135(6):2339–2354, 2007.

[33] R. Giering and T. Kaminski. Recipes for adjoint code construction. *ACM Trans. On Math. Software*, 24:437–474, 1998.

[34] S. Gratton, A.S. Lawless, and N.K. Nichols. Approximate gauss-newton methods for nonlinear least squares problems. *SIAM J. Optim.*, 18(1):106–132, 2007.

[35] A.K. Griffith and N.K Nichols. Adjoint methods in data assimilation for estimating model error. *Flow, Turbulence and Combustion*, 65:469–488, 2000.

[36] V. Guidard and C. Fischer. Introducing the coupling information in a limited-area variational assimilation. *Quart. J. Roy. Meteor. Soc.*, 134(632):723–735, 2008.

[37] N. Gustafsson, L. Berre, S. Hörnquist, X.Y. Huang, M. Lindskog, B. Navascués, K.S. Mogensen, and S. Thorsteinsson. Three-dimensional variational data assimilation for a limited area model. *Tellus A*, 53(4):425–446, 2001.

[38] N. Gustafsson, X.Y. Huang, X. Yang, K. Mogensen, M. Lindskog, O. Vignes, T. Wilhelmsson, and S. Thorsteinsson. Four-dimensional variational data assimilation for a limited area model. *Tellus A*, 64:DOI: 10.3402/tellusa.v64i0.14985, 2012.

[39] S.A. Haben. *Conditioning and preconditioning of the minimisation problem in variational data assimilation.* PhD thesis, Department of Mathematics and Statistics, University of Reading, 2011.

[40] S.A. Haben, A.S. Lawless, and N.K. Nichols. Conditioning and preconditioning of the variational data assimilation problem. *Computers & Fluids*, 46(1):252–256, 2011.

[41] S.A. Haben, A.S. Lawless, and N.K. Nichols. Conditioning of incremental variational data assimilation, with application to the met office system. *Tellus A*, 63:782–792, 2011.

[42] B.A. Harris and G. Kelly. A satellite radiance-bias correction scheme for data assimilation. *Quart. J. Roy. Meteor. Soc.*, 127(574):1453–1468, 2001.

[43] S.B. Healy and A.A. White. Use of discrete fourier transforms in the 1d-var retrieval problem. *Quart. J. Roy. Meteor. Soc.*, 131(605):63–72, 2005.

[44] F. Hilton, N.C. Atkinson, S.J. English, and J.R. Eyre. Assimilation of iasi at the met office and assessment of its impact through observing system experiments. *Quart. J. Roy. Meteor. Soc.*, 135(639):495–505, 2009.

[45] A. Hollingsworth and P. Lönnberg. The statistical structure of short-range forecast errors as determined from radiosonde data. part i: The wind field. *Tellus A*, 38(2):111–136, 1986.

[46] B.J. Hoskins, M.E. McIntyre, and A.W. Robertson. On the use and significance of isentropic potential vorticity maps. *Quart. J. Roy. Meteor. Soc.*, 111:877–946, 1995.

[47] I. Hoteit and A. Köhl. Efficiency of reduced-order, time-dependent adjoint data assimilation approaches. *Journal of oceanography*, 62(4):539–550, 2006.

[48] X.Y. Huang, Q. Xiao, D.M. Barker, X. Zhang, J. Michalakes, W. Huang, T. Henderson, J. Bray, Y. Chen, Z. Ma, J. Dudhia, Y. Guo, X. Zhang, D-J. Won, H-C. Lin, and Y-H. Kuo. Four-dimensional variational data assimilation for wrf: Formulation and preliminary results. *Mon. Wea. Rev.*, 137(1):299–314, 2009.

[49] N.B. Ingleby. The statistical structure of forecast errors and its representation in the Met Office global 3-dimensional variational data assimilation scheme. *Quart. J. Roy. Meteor. Soc.*, 127(571):209–231, 2001.

[50] N.B. Ingleby and A.C. Lorenc. Bayesian quality control using multivariate normal distributions. *Quart. J. Roy. Meteor. Soc.*, 119(513):1195–1225, 1993.

[51] M. Janiskova, J-N. Thepaut, and J-F. Geleyn. Simplified and regular physical parametrizations for incremental four-dimensional variational assimilation. *Mon. Wea. Rev.*, 127:26–45, 1999.

[52] C. Johnson, B.J. Hoskins, and N.K. Nichols. A singular vector perspective of 4d-var: Filtering and interpolation. *Quart. J. Roy. Meteor. Soc.*, 131(605):1–19, 2005.

[53] D. Katz, A.S. Lawless, N.K. Nichols, M.J.P. Cullen, and R.N. Bannister. Correlations of control variables in variational data assimilation. *Quart. J. Roy. Meteor. Soc.*, 137:620–630, 2011.

[54] T. Kawabata, H. Seko, K. Saito, T. Kuroda, K. Tamiya, T. Tsuyuki, Y. Honda, and Y. Wakazuki. An assimilation and forecasting experiment of the nerima heavy rainfall with a cloud-resolving nonhydrostatic 4-dimensional variational data assimilation system. *J. Met. Soc. Japan. Ser. II*, 85(3):255–276, 2007.

[55] S. Laroche and P. Gauthier. A validation of the incremental formulation of 4D variational data assimilation in a nonlinear barotropic flow. *Tellus*, 50A:557–572, 1998.

[56] A.S. Lawless. A note on the analysis error associated with 3d-fgat. *Quart. J. Roy. Meteor. Soc.*, 136(649):1094–1098, 2010.

[57] A.S. Lawless, S. Gratton, and N.K. Nichols. Approximate iterative methods for variational data assimilation. *International journal for numerical methods in fluids*, 47(10-11):1129–1135, 2005.

[58] A.S. Lawless, S. Gratton, and N.K. Nichols. An investigation of incremental 4d-var using non-tangent linear models. *Quart. J. Roy. Meteor. Soc.*, 131(606):459–476, 2005.

[59] A.S. Lawless and N.K. Nichols. Inner-loop stopping criteria for incremental four-dimensional variational data assimilation. *Mon. Wea. Rev.*, 134(11):3425–3435, 2006.

[60] A.S. Lawless, N.K. Nichols, and S.P. Ballard. A comparison of two methods for developing the linearization of a shallow-water model. *Quart. J. Roy. Meteor. Soc.*, 129:1237–1254, 2003.

[61] A.S. Lawless, N.K. Nichols, C. Boess, and A. Bunse-Gerstner. Approximate gauss–newton methods for optimal state estimation using reduced-order models. *International journal for numerical methods in fluids*, 56(8):1367–1373, 2008.

[62] A.S. Lawless, N.K. Nichols, C. Boess, and A. Bunse-Gerstner. Using model reduction methods within incremental four-dimensional variational data assimilation. *Mon. Wea. Rev.*, 136(4):1511–1522, 2008.

[63] F.-X. Le Dimet and O. Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus*, 38A:97–110, 1986.

[64] D.J. Lea, J.P. Drecourt, K. Haines, and M.J. Martin. Ocean altimeter assimilation with observational- and model-bias correction. *Quart. J. Roy. Meteor. Soc.*, 134(636):1761–1774, 2008.

[65] Y. Li, I.M. Navon, W. Yang, X. Zou, J.R. Bates, S. Moorthi, and R.W. Higgins. Four-dimensional variational data assimilation experiments with a multilevel semi-Lagrangian semi-implicit general circulation model. *Mon. Wea. Rev.*, 122:966–983, 1994.

[66] M. Lindskog, D. Dee, Y. Trémolet, E. Andersson, G. Radnóti, and M. Fisher. A weak-constraint four-dimensional variational analysis system in the stratosphere. *Quart. J. Roy. Meteor. Soc.*, 135(640):695–706, 2009.

[67] C. Liu, Q. Xiao, and B. Wang. An ensemble-based four-dimensional variational data assimilation scheme. part i: Technical formulation and preliminary test. *Mon. Wea. Rev.*, 136(9):3363–3373, 2008.

[68] A.C. Lorenc. Analysis methods for numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, 112:1177–1194, 1986.

[69] A.C. Lorenc. Development of an operational variational assimilation scheme. *J. Met. Soc. Japan*, 75:339–346, 1997.

[70] 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, and D. Li. The met. office global 3-dimensional variational data assimilation scheme. *Quart. J. Roy. Meteor. Soc*, 126:2991–3012, 2000.

[71] A.C. Lorenc, R.B. Bell, and B. Macpherson. The Meteorological Office analysis correction data assimilation scheme. *Quart. J. Roy. Meteor. Soc.*, 117:59–89, 1991.

[72] A.C. Lorenc and T. Payne. 4d-var and the butterfly effect: Statistical four-dimensional data assimilation for a wide range of scales. *Quart. J. Roy. Meteor. Soc.*, 133(624):607–614, 2007.

[73] A.C. Lorenc and F. Rawlins. Why does 4d-var beat 3d-var? *Quart. J. Roy. Meteor. Soc.*, 131(613):3247–3257, 2005.

[74] J.F. Mahfouf and F. Rabier. The ecmwf operational implementation of four-dimensional variational assimilation. ii: Experimental results with improved physics. *Quart. J. Roy. Meteor. Soc.*, 126(564):1171–1190, 2000.

[75] T. Mochizuki, N. Sugiura, T. Awaji, and T. Toyoda. Seasonal climate modeling over the indian ocean by employing a 4d-var coupled data assimilation approach. *Journal of Geophysical Research*, 114(C11):C11003, 2009.

[76] A.M. Moore, H.G. Arango, G. Broquet, B.S. Powell, A.T. Weaver, and J. Zavala-Garay. The regional ocean modeling system (roms) 4-dimensional variational data assimilation systems part i–system overview and formulation. *Progress in Oceanography*, 91:34–49, 2011.

[77] I.M. Navon, X. Zou, J. Derber, and J. Sela. Variational data assimilation with an adiabatic version of the NMC spectral model. *Mon. Wea. Rev.*, 120:1433–1446, 1992.

[78] J. Nocedal and S.J. Wright. *Numerical optimization*. Springer Verlag, 2006.

[79] S.K. Park and D. Zupanski. Four-dimensional variational data assimilation for mesoscale and storm-scale applications. *Meteorology and Atmospheric Physics*, 82(1):173–208, 2003.

[80] D.F. Parrish and J.C. Derber. The National Meteorological Center's spectral statistical-interpolation analysis system. *Mon. Wea. Rev.*, 120:1747–1763, 1992.

[81] F. Rabier and P. Courtier. Four-dimensional assimilation in the presence of baroclinic instability. *Quart. J. Roy. Meteor. Soc.*, 118:649–672, 1992.

[82] F. Rabier, H. Järvinen, E. Klinker, J.F. Mahfouf, and A. Simmons. The ecmwf operational implementation of four-dimensional variational assimilation. i: Experimental results with simplified physics. *Quart. J. Roy. Meteor. Soc.*, 126(564):1143–1170, 2000.

[83] F. Rawlins, S.P. Ballard, K.J. Bovis, A.M. Clayton, D. Li, G.W. Inverarity, A.C. Lorenc, and T.J. Payne. The met office global four-dimensional variational data assimilation scheme. *Quart. J. Roy. Meteor. Soc.*, 133(623):347–362, 2007.

[84] C. Robert, S. Durbiano, E. Blayo, J. Verron, J. Blum, and F.X. Le Dimet. A reduced-order strategy for 4d-var data assimilation. *Journal of Marine Systems*, 57(1):70–82, 2005.

[85] N. Rostaing, S. Dalmas, and A. Galligo. Automatic differentiation in Odyssée. *Tellus*, 45A:558–568, 1993.

[86] Y. Sasaki. An objective analysis based on the variational method. *J. Met. Soc. Japan*, 36:77–88, 1958.

[87] Y. Sasaki. Some basic formalisms in numerical variational analysis. *Mon. Wea. Rev.*, 98:875–883, 1970.

[88] L.M. Stewart. *Correlated observation errors in data assimilation.* PhD thesis, Department of Mathematics, University of Reading, 2010.

[89] N. Sugiura, T. Awaji, S. Masuda, T. Mochizuki, T. Toyoda, T. Miyama, H. Igarashi, and Y. Ishikawa. Development of a four-dimensional variational coupled data assimilation system for enhanced analysis and prediction of seasonal to interannual climate variations. *Journal of Geophysical Research*, 113(C10):C10017, 2008.

[90] O. Talagrand and P. Courtier. Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory. *Quart. J. Roy. Meteor. Soc.*, 113:1311–1328, 1987.

[91] J-N. Thepaut and P. Courtier. Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model. *Quart. J. Roy. Meteor. Soc.*, 117:1225–1254, 1991.

[92] Y. Trémolet. Accounting for an imperfect model in 4d-var. *Quart. J. Roy. Meteor. Soc.*, 132(621):2483–2504, 2006.

[93] Y. Trémolet. Model-error estimation in 4d-var. *Quart. J. Roy. Meteor. Soc.*, 133(626):1267–1280, 2007.

[94] J. Tshimanga, S. Gratton, A.T. Weaver, and A. Sartenaer. Limited-memory preconditioners, with application to incremental four-dimensional variational data assimilation. *Quart. J. Roy. Meteor. Soc.*, 134(632):751–769, 2008.

[95] J. Vialard, A.T. Weaver, D.L.T. Anderson, and P. Delecluse. Three-and four-dimensional variational assimilation with a general circulation model of the tropical pacific ocean. part ii: Physical validation. *Mon. Wea. Rev.*, 131(7):1379–1395, 2003.

[96] A. Weaver and P. Courtier. Correlation modelling on the sphere using a generalized diffusion equation. *Quart. J. Roy. Meteor. Soc.*, 127(575):1815–1846, 2001.

[97] A.T. Weaver, C. Deltel, E. Machu, S. Ricci, and N. Daget. A multivariate balance operator for variational ocean data assimilation. *Quart. J. Roy. Meteor. Soc.*, 131(613):3605–3625, 2005.

[98] A.T. Weaver and I. Mirouze. On the diffusion equation and its application to isotropic and anisotropic correlation modelling in variational assimilation. *Quart. J. Roy. Meteor. Soc.*, page doi: 10.1002/qj.1955, 2012.

[99] A.T. Weaver, J. Vialard, and D.L.T. Anderson. Three-and four-dimensional variational assimilation with a general circulation model of the tropical pacific ocean. part i: Formulation, internal diagnostics, and consistency checks. *Mon. Wea. Rev.*, 131(7):1360–1378, 2003.

[100] Q. Xu. Generalized adjoint for physical processes with parameterized discontinuities. Part I: Basic issues and heuristic examples. *J. Atmos. Sci.*, 53:1123–1155, 1996.

[101] D. Zupanski. A general weak constraint applicable to operational 4DVAR data assimilation systems. *Mon. Wea. Rev.*, 125:2274–2292, 1997.