RMetS
Royal Meteorological Society

Impact of non-smooth observation operators on variational and sequential data assimilation for a limited-area shallow-water equation model

J. L. Steward,[a]* I. M. Navon,[a] M. Zupanski[b] and N. Karmitsa[c]
[a]*Department of Scientific Computing, The Florida State University, Tallahassee, USA*
[b]*Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, USA*
[c]*Department of Mathematics, University of Turku, Turku, Finland*

*Correspondence to: J. L. Steward, The Florida State University, 400 Dirac Science Library, Tallahassee, FL 32306-4120, USA. E-mail: jls07c@fsu.edu

We investigate the issue of variational and sequential data assimilation with nonlinear and non-smooth observation operators using a two-dimensional limited-area shallow-water equation model and its adjoint. The performance of the four-dimensional variational approach (4D-Var: two dimensions plus time) compared with that of the maximum-likelihood ensemble filter (MLEF), a hybrid ensemble/variational method, is tested in the presence of non-smooth observation operators.

Following the work of Lewis & Overton and Karmitsa, we investigate minimization of the data-assimilation cost functional using the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) quasi-Newton algorithm originally intended for smooth optimization and the limited-memory bundle method (LMBM) algorithm specifically designed to address large-scale non-smooth minimization problems.

Numerical results obtained for the MLEF method show that the LMBM algorithm yields results superior to the L-BFGS method. Results for 4D-Var suggest that L-BFGS performs well when the non-smoothness is not extreme, but fails for non-smooth functions with large Lipschitz constants. The LMBM method is found to be a suitable choice for large-scale non-smooth optimization, although additional work is needed to improve its numerical stability. Finally, the results and methodologies of 4D-Var and MLEF are compared and contrasted. Copyright © 2011 Royal Meteorological Society

Key Words: 4D-Var; MLEF; L-BFGS; LMBM; non-smooth optimization

Received 6 March 2011; Revised 14 June 2011; Accepted 26 August 2011; Published online in Wiley Online Library 4 October 2011

Citation: Steward JL, Navon IM, Zupanski M, Karmitsa N. 2012. Impact of non-smooth observation operators on variational and sequential data assimilation for a limited-area shallow-water equation model. *Q. J. R. Meteorol. Soc.* **138**: 323–339. DOI:10.1002/qj.935

1. Introduction

In this article, we investigate data assimilation of the shallow-water equations with discontinuous observation operators in order to compare the performance of large-scale non-smooth optimization methods. We employ both a variational (4D-Var: Chevallier *et al.*, 2004) approach and the maximum-likelihood ensemble filter (MLEF: Zupanski, 2005; Zupanski and Zupanski, 2006) hybrid ensemble/variational data-assimilation method. In

light of Zupanski *et al.* (2008)'s demonstration that MLEF can be derived without a differentiability requirement for the prediction model and the observation operators, we investigate the non-smooth optimization properties of MLEF. We also compare and contrast the results obtained using the limited-memory Broyden–Fletcher–Goldfarb–Shanno method (L-BFGS: Liu and Nocedal, 1989) and limited-memory bundle method (LMBM: Karmitsa, 2007) minimization algorithms for large-scale non-smooth optimization within 4D-Var and MLEF.

Data assimilation (Kalnay, 2003) aims to utilize observations of a system optimally in combination with a previous estimate of the initial state of the system known as the *background*. Both the background and observations contain noise and thus must be considered together. An observation operator is a function mapping between model space and the observations. For example, in the context of atmospheric science, an observation operator for satellite radiances has a radiative transfer model that can take the state of the atmosphere as input and produce synthetic observations that can be compared with actual satellite radiance observations. Non-smooth observation operators are those that have a discontinuity in the derivative of order zero (function value) or higher. The lower the order of the discontinuity, the more difficult the issue is to deal with numerically and mathematically (Makela and Neittaanmaki, 1992). In this work we focus our investigation on observation operators with discontinuities in the first derivative.

The issue of data assimilation with discontinuous observation operators is relevant to many outstanding data-assimilation problems. For example, the data assimilation of 'all-sky' satellite radiance observations, which may or may not be affected by clouds, has a discontinuous observation operator with respect to cloud microphysical variables (Zou and Navon, 1996; Janiskova *et al.*, 2002; Errico *et al.*, 2007). Sea-ice modelling is another area in which discontinuities complicate the data-assimilation process (Levy *et al.*, 2010). Furthermore, in many other situations discontinuities arise due to 'on/off' switches in model parametrizations (Zupanski, 1993; Zupanski and Mesinger, 1995; Zhang *et al.*, 2001; Errico *et al.*, 2007). These discontinuities may become an issue for any smooth optimization algorithm that is employed to minimize the cost functional in either the 4D-Var system or hybrid ensemble/variational methods such as MLEF. Their presence also poses serious problems for the correct adjoint-model formulation. The non-smoothness may lead to a poor solution or even a failure of the minimization algorithm (Greenwald *et al.*, 2004). In particular, issues may arise with both the line search and descent direction (Makela and Neittaanmaki, 1992). This article evaluates several optimization algorithms in data assimilation using a finite-difference shallow-water equation model.

Because a typical data-assimilation analysis equation proceeds from the assumption of linearity and smoothness, current approaches include regularizing and smoothing simplified versions of the parametrizations (van Leeuwen, 2001; Janiskova and Morcrette, 2005). This approach can introduce new problems, however (Zhang *et al.*, 2000). In addition, when the discontinuity in the observation process is mathematically inherent with respect to the control variables –as is the case for all-sky radiances when cloud microphysical variables are part of the control –it is not clear that smoothing or regularizing the problem is the correct approach. Changing the problem being solved may indeed prevent the extraction of the maximum amount of information from highly nonlinear or discontinuous processes. This limitation would in turn affect the skill of the data-assimilation procedure.

MLEF has been tested with the shallow-water equations in Zupanski *et al.* (2006), Uzunoglu *et al.* (2007) and Fletcher and Zupanski (2008) and non-smooth variational data assimilation problems have been investigated in Makela and Neittaanmaki (1992), Homescu and Navon (2003), Zhang *et al.* (2000), Zhu *et al.* (2002) and Bardos and Pironneau (2005) for highly simplified problems. However, the existing optimization algorithms at the time those studies were conducted were not suitable for large-scale non-convex optimization (Haarala *et al.*, 2004). The situation has now changed, due to both theoretical and algorithmic advances. An excellent comparison of non-smooth optimization algorithms for large-scale minimization is given in Karmitsa *et al.* (2009), showing positive results for problems with as many as 4000 variables. Levy and coauthors recently investigated a physical-based approach for potentially discontinuous optimal interpolation of sea-ice data assimilation for a model with 15 physical control variables in Levy *et al.* (2010). The approach in the current article suggests techniques suitable for more general and larger-scale data-assimilation problems, although combining these techniques is certainly an intriguing possibility.

The smooth optimization quasi-Newton L-BFGS algorithm, long used in data assimilation (Zou *et al.*, 1992; Honda *et al.*, 2005), has recently been found to possess the properties of a non-smooth optimization algorithm in Lewis and Overton (2008a,b) and Skajaa (2010). This method may offer promise for large-scale non-smooth optimal control problems.

In this article we test the L-BFGS algorithm as well as the LMBM algorithm, a globally convergent non-smooth optimization algorithm specifically designed for large-scale, possibly non-convex minimization (Haarala *et al.*, 2007; Karmitsa, 2007) for problems with more than 1000 control variables. In addition, we also test the MLEF hybrid-filter data-assimilation method implemented with the L-BFGS minimization algorithm to assess its non-smooth optimization properties (Zupanski *et al.*, 2008).

We conduct our tests using a series of closely related 4D-Var optimal control problems with observation operators containing varying degrees of non-smoothness. We also test the same data-assimilation problems using the hybrid MLEF filter.

The article is organized as follows: section 2 describes the nonlinear limited-area shallow-water equation model and the initial and boundary conditions employed, section 3 details the data-assimilation problem and methodologies used to solve it, section 4 describes the non-smooth optimization algorithms used in this study, section 5 details the set-up of the computational experiments and section 6 presents the numerical results. Finally, section 7 is reserved for a summary and conclusions.

2. Shallow-water equation model

We begin by describing the model that will be used in our optimal control problem.

Consider the limited-area shallow-water equation model as detailed in Wang *et al.* (1992):

$$\frac{\partial u}{\partial t} = -u\frac{\partial u}{\partial x} - v\frac{\partial u}{\partial y} + fv - \frac{\partial \phi}{\partial x}, \tag{1}$$

$$\frac{\partial v}{\partial t} = -u\frac{\partial v}{\partial x} - v\frac{\partial v}{\partial y} - fu - \frac{\partial \phi}{\partial y}, \tag{2}$$

$$\frac{\partial \phi}{\partial t} = -\frac{\partial u\phi}{\partial x} - \frac{\partial v\phi}{\partial y}, \tag{3}$$

where u and v are the two components of the horizontal velocity in m s^{-1}, $\phi = gh$ is the geopotential field in m^2 s^{-2}, h is the free surface height in m and f is the Coriolis factor in s^{-1}.

The initial conditions used were based on those in Grammeltvedt (1969), namely a channel on a β plane of length L and depth D, with h given by

$$\begin{aligned} h(x,y) &= h_0 + h_1 \tanh\left(\frac{9(y-y_0)}{2D}\right) \\ &\quad + h_2 \operatorname{sech}^2\left(\frac{9(y-y_0)}{2D}\right)\sin\left(\frac{2\pi x}{L}\right), \end{aligned} \tag{4}$$

where $h_0 = 2000$ m, $h_1 = -220$ m, $h_2 = 133$ m, $L = 6000$ km, $D = 4400$ km and $y_0 = D/2$.

From Eq. (4), the initial conditions are derived through geostrophic balance by the relation

$$\begin{aligned} \phi_0(x,y) &= gh(x,y), \\ u_0(x,y) &= -\frac{g}{f}\frac{\partial h}{\partial y}(x,y), \\ v_0(x,y) &= \frac{g}{f}\frac{\partial h}{\partial x}(x,y), \end{aligned} \tag{5}$$

where $g = 10$ m s^{-2} and $f = 10^{-4}$ s^{-1}.

This model is discretized using the second-order quadratic conservation advective scheme detailed in Grammeltvedt (1969), referred to as 'scheme F'. The space and time increments are $\Delta x = 300$ km, $\Delta y = 220$ km and $\Delta t = 600$ s, respectively, resulting in a mesh comprising 21 × 21 spatial grid points. The model is integrated for 80 time steps, i.e. a window of assimilation of 13 h 20 min in model time. These initial conditions are shown in Figure 1.

Because in this article the u and v velocity components will have separate observation-operator components, as described in section 5.2, the contour plots for the initial values of these two fields are shown in Figure 2.

The boundary conditions are given by a rigid-wall homogeneous Neumann condition along the south and north boundaries and wrapping periodic conditions along the east/west boundary. In other words,

$$\begin{aligned} u(x_\mathrm{l}, y, t) &= u(x_\mathrm{r}, y, t), \\ v(x_\mathrm{l}, y, t) &= v(x_\mathrm{r}, y, t), \\ \phi(x_\mathrm{l}, y, t) &= \phi(x_\mathrm{r}, y, t), \frac{\partial u}{\partial y}(x, y_\mathrm{t}, t) = 0, \\ \frac{\partial u}{\partial y}(x, y_\mathrm{b}, t) &= 0, \end{aligned}$$

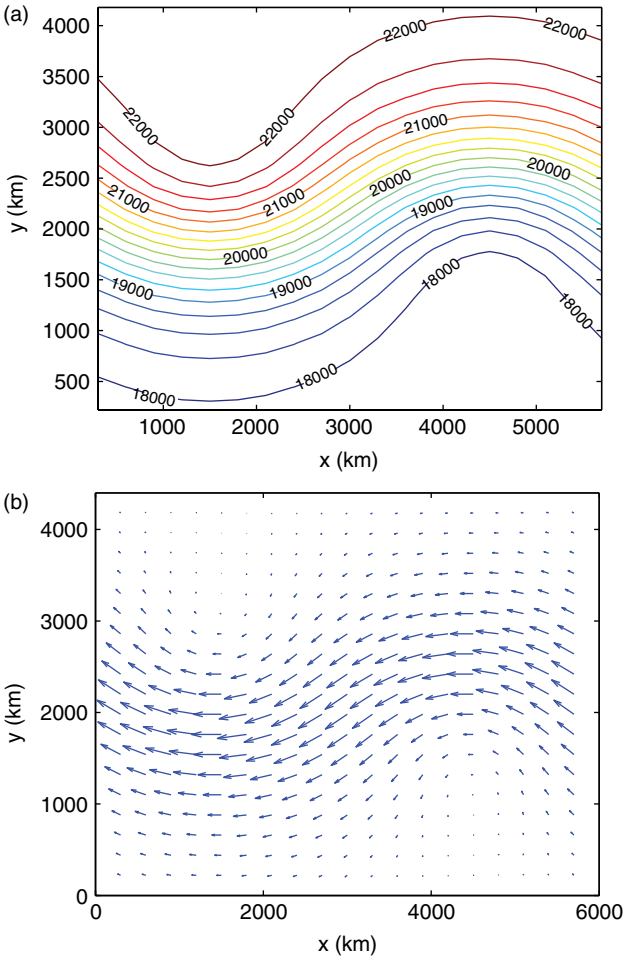

Figure 1. Initial conditions (u_0, v_0, ϕ_0): (a) ϕ_0 from (5) and (b) wind field from (5). Arrows are scaled by a factor of 1 km. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

$$\begin{aligned} v(x, y_\mathrm{t}, t) &= 0, \\ v(x, y_\mathrm{b}, t) &= 0, \\ \frac{\partial \phi}{\partial y}(x, y_\mathrm{t}, t) &= 0, \\ \frac{\partial \phi}{\partial y}(x, y_\mathrm{b}, t) &= 0, \end{aligned}$$

where x_l, x_r, y_t, y_b are locations of the left, right, top and bottom boundaries, respectively.

3. Data-assimilation methods

In a general setting with standard assumptions, the data-assimilation problem is shown in Table I.

The background state x_b is assumed to be of the form

$$x_\mathrm{b} = x_\mathrm{true}^{(0)} + \eta_\mathrm{b}, \tag{6}$$

where x_true is the true state of the system and $\eta_\mathrm{b} \sim N(0, B)$ is a Gaussian random variable with mean 0 and covariance matrix B.

The observations $y^{(i)}$ are assumed to be

$$y^{(i)} = \mathcal{H}\left(x_\mathrm{true}^{(i)}\right) + \eta_\mathrm{obs}, \tag{7}$$

where $\eta_\mathrm{obs} \sim N(0, R)$ is a Gaussian random variable with mean 0 and covariance matrix R. The observation errors are considered independent, so R is a diagonal matrix.

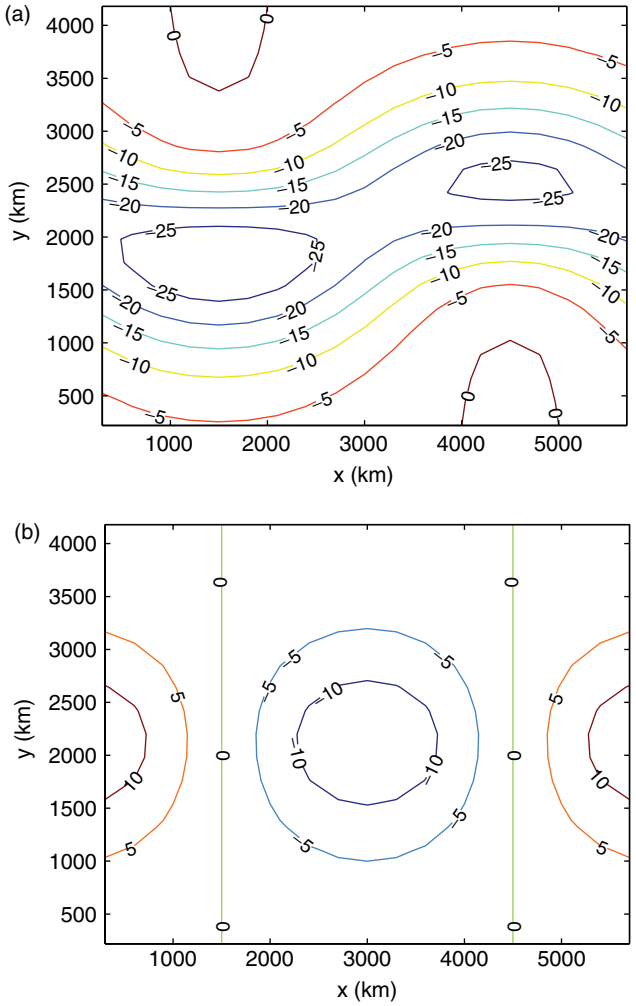

Figure 2. Contour plot for (a) u_0 and (b) v_0. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

The model $\mathcal{M}$ we investigate here consists of the shallow-water equations detailed in section 2. As we will show below in section 5, we create a twin experiment by taking the initial conditions in Eq. (5) to be $x^{(0)}_{\text{true}}$ in order to create the background and observations.

Many different algorithms have been employed to solve the data-assimilation problem in Table I. In this article, we investigate the solution of the data-assimilation problem using both 4D-Var, a variational approach, and the MLEF, which employs both variational and sequential techniques and can thus be considered a hybrid method.

In practice there are many issues that arise when trying to solve the data-assimilation problem. Issues such as estimating the uncertainties B and R, observation sparsity, model error and background-state formulation are not considered in this research in order to focus on the case of an observation operator $\mathcal{H}$ that is non-differentiable (as detailed in section 5.2).

3.1. 4D-Var

The 4D-Var approach to data assimilation is to minimize a cost function, with the lack of fit between background and observations scaled by their respective uncertainties. In other words, we seek the value x^* that minimizes

$$J(x) = \tfrac{1}{2}\,\delta_{\text{b}}(x)^{\text{T}} B^{-1}\delta_{\text{b}}(x) + \tfrac{1}{2}\sum_{k=0}^{NT}\delta_{y_k}(x)^{\text{T}} R^{-1}\delta_{y_k}(x), \tag{8}$$

subject to the constraint

$$x_i(x) = \mathcal{M}\left[x_{i-1}(x)\right], \qquad x_0(x) = x, \tag{9}$$

where x is the control variable representing the unknown initial conditions,

$$\delta_{\text{b}}(x) = x - x_{\text{b}} \tag{10}$$

is the difference between the control variable and the background state x_{b} and

$$\delta_{y_k}(x) = y^{(k)} - \mathcal{H}(x^{(k)}(x)) \tag{11}$$

is the difference between the observation at time k, $y^{(k)}$, and the model at time k, $x^{(k)}(x)$, acted on by the observation operator $\mathcal{H}$.

In this work, the model $\mathcal{M}$ is taken as a strong constraint; in other words, the model constraint equation (9) is satisfied at each time step. Once the optimal initial conditions x^* have been found, the optimal trajectory $x^{(i)}(x^*)$ can be found by evolving the model forward in time using model $\mathcal{M}$.

In order to proceed with the minimization, the gradient of the cost function (8) is needed. In the case of a once-differentiable model $\mathcal{M}$ and operator $\mathcal{H}$, this is given by

$$\nabla_x J(x) = B^{-1}\delta_{\text{b}} - \sum_{k=0}^{NT}\frac{\partial x^{(k)}}{\partial x}\frac{\partial \mathcal{H}}{\partial x^{(k)}}R^{-1}\delta_{y_k}, \tag{12}$$

where $\partial\mathcal{H}/\partial x^i$ is the Jacobian of the observation operator and $\partial x^{(i)}/\partial x \in \mathbb{R}^{3NM\times 3NM}$ is the Jacobian of $x^{(i)}$. From Eq. (9),

$$\frac{\partial x^{(i)}}{\partial x}(x) = \frac{\partial x^{(i-1)}}{\partial x}(x)\mathbf{M}\left(x^{(i-1)}(x)\right) \tag{13}$$

for $i = 1, \ldots, NT$, where $\mathbf{M}$ is the Jacobian of the model operator.

In order to compute this gradient efficiently, the adjoint of the tangent linear model is computed. The adjoint is detailed in Wang *et al.* (1992). It uses a discretize-then-differentiate approach (Gunzburger, 2003), i.e. the derivatives are found by applying the chain rule to the discrete operator $\mathcal{M}$ and reversing the order of the Fortran code.

As detailed in sections 4 and 5, the functions $\mathcal{H}$ in this research are only assumed to be locally Lipschitz-continuous and may not in general be differentiable (i.e. they are non-smooth).

3.2. MLEF

The MLEF approach to data assimilation shares similarities with both a variational approach such as 4D-Var and an ensemble approach such as the Ensemble Kalman Filter (see e.g. Kalnay *et al.*, 2007) and especially the Ensemble Square Root Filter (Tippett *et al.*, 2003) and Ensemble Kalman Transform Filter (Bishop *et al.*, 2001). The derivation presented below is similar to those of Fletcher and Zupanski

Table I. Data assimilation problem. The following information is required to find an 'optimal' trajectory $x^{(i)} \in \mathbb{R}^{N_{\text{state}}}$ at each time step $\{i = 0, \ldots, NT\}$.

A 'background' state of the model x_{b} that approximates $x^{(0)}$ (from either a first guess or a previous prediction).
Complete or partial noisy observations of the system $y^{(j)} \in \mathbb{R}^{N_{\text{obs}}^{(j)}}$ for some or all values $0 \leq j \leq NT$.
Background-error covariance matrix $B \in \mathbb{R}^{N_{\text{state}} \times N_{\text{state}}}$ quantifying the covariance of the error between the background x_{b} and the unknown true state.
Observation-error covariance matrix $R_j \in \mathbb{R}^{N_{\text{obs}}^{(j)} \times N_{\text{obs}}^{(j)}}$ quantifying the covariance of the error between the observations $y^{(j)}$ and the unknown true observations.
A model $\mathcal{M} : \mathbb{R}^{N_{\text{state}}} \rightarrow \mathbb{R}^{N_{\text{state}}}$ that maps $x^{(i)}$ to $x^{(i+1)}$ (assumed here to be perfect).
An observation operator $\mathcal{H}_j : \mathbb{R}^{N_{\text{state}}} \rightarrow \mathbb{R}^{N_{\text{obs}}^{(j)}}$ (also assumed to be perfect) that models the mapping between $x^{(j)}$ to the observations $y^{(j)}$.

(2008) and Zupanski *et al.* (2008), in which additional details can be found.

Like the Kalman filter families, the MLEF algorithm proceeds in two stages: forecast and analysis. Suppose that a square-root analysis covariance matrix $\left(P_{\text{a}}^{(k-1)}\right)^{1/2}$ is available at time $k-1$, such that

$$P_{\text{a}}^{(k-1)} = \left(P_{\text{a}}^{(k-1)}\right)^{1/2} \left(P_{\text{a}}^{(k-1)}\right)^{T/2}. \tag{14}$$

This is approximately equal to the background B at time 0, i.e. $P_{\text{a}}^{(0)} \approx B$. In MLEF, the columns of the forecast-error covariance $\left(P_{\text{f}}^{(k)}\right)^{1/2}$ at time k are given by

$$\begin{aligned} \left(P_{\text{f}}^{(k)}\right)^{1/2} &= \left[p_1^{\text{f}}, p_2^{\text{f}}, \ldots, p_{N_{\text{ens}}}^{\text{f}}\right], \\ p_i^{\text{f}} &= \mathcal{M}\left(x_{\text{a}}^{(k-1)} + p_i^{\text{a}}\right) - \mathcal{M}\left(x_{\text{a}}^{(k-1)}\right), \end{aligned} \tag{15}$$

where N_{ens} are the number of ensembles, $x_{\text{a}}^{(k-1)}$ is the previous analysis value ($x_{\text{a}}^{(0)} = x_{\text{b}}$) and p_i^a is the ith column of $\left(P_{\text{a}}^{(k-1)}\right)^{1/2}$.

Once the forecast covariance has been obtained, the analysis step can proceed. The analysis step of MLEF takes its inspiration from variational methods, in particular 3D-Var (Lorenc, 1986). MLEF seeks to find the analysis, $x_{\text{a}}^{(k)} = \arg\min_x J$, that minimizes the cost function

$$J(x) = \frac{1}{2}\, \delta_{\text{f}}(x)^{\text{T}}\, P_{\text{f}}^{-1} \delta_{\text{f}}(x) + \frac{1}{2}\, \delta_{y_k}(x)^{\text{T}}\, R^{-1} \delta_{y_k}(x), \tag{16}$$

where $\delta_{\text{f}} = x - x_{\text{f}}^{(k)}$, $x_{\text{f}}^{(k)} = \mathcal{M}\left(x_{\text{a}}^{(k-1)}\right)$ and the other variables are the same as in section 3.1.

In order to avoid inverting the matrix P_{f}, which will be rank-deficient as it is approximated by a matrix of rank at most N_{ens}, a change of variables is introduced, namely

$$\delta_{\text{f}}(x) = P_{\text{f}}^{1/2}\, (I + C(x))^{-\text{T}/2}\, \zeta, \tag{17}$$

where ζ are the new control variables, defined in the $N_{\text{ens}} \times 1$ space of ensembles, and $C \in \mathbb{R}^{N_{\text{ens}} \times N_{\text{ens}}}$ is a preconditioning matrix of the quadratic cost function (16). C is formed as follows:

$$C(x) = Z^{\text{T}} Z, \tag{18}$$

where

$$Z_i(x) = R^{-1/2} \left[\mathcal{H}(x + p_i^{\text{f}}) - \mathcal{H}(x)\right]. \tag{19}$$

Because of the mutual dependence between the preconditioner and x, C is fixed to $C(x = x_{\text{f}}^{(k)})$, i.e. C does not change during the minimization process.

The inversion of the symmetric matrix $(I + C)$, required by (17), is accomplished using a spectral decomposition of the form

$$I + C = V \Lambda V^{\text{T}}, \tag{20}$$

where V is an orthogonal matrix of eigenvectors and Λ is a diagonal matrix of eigenvalues. Once this decomposition is found, the required square root can be found by

$$(I + C)^{-1/2} = V \Lambda^{-1/2} V^{\text{T}}. \tag{21}$$

Finally, the square-root analysis covariance matrix at time (k) is found by

$$P_{\text{a}}^{1/2} = P_{\text{f}}^{1/2} \left[I + C(x_{\text{a}}^{(k)})\right]^{-\text{T}/2}, \tag{22}$$

with the notation $C(x_{\text{a}}^{(k)})$ denoting that C is recomputed at the solution $x_{\text{a}}^{(k)}$ of (16).

This update to the covariance matrix is similar to that of Bishop *et al.* (2001); however, the main difference is that the observation operator is not restricted to be linear and the Jacobian is not required.

In summary, MLEF is an ensemble method that directly maximizes the posterior probability density function at each time step. It does not require the Jacobian of either the model or the observation operator.

4. Non-smooth optimization algorithms

4.1. Introduction

In this section we describe the non-smooth optimization algorithms used in this research. Before further discussion we introduce a few common definitions. In what follows we use an Euclidean norm, i.e. $||x|| = \left(\sum_{i=1}^{n} x_i^2\right)^{1/2}$.

Definition 4.1 *Function* $J : \mathbb{R}^n \rightarrow \mathbb{R}$ *is* locally Lipschitz-continuous *at* $x \in \mathbb{R}^n$ *with a constant* $L > 0$ *if there exists a positive number* ϵ *such that*

$$|f(y) - f(z)| \leq L||y - z||$$

for all y, z *such that* $||x - y|| \leq \epsilon$, $||x - z|| \leq \epsilon$.

Intuitively, L is an upper limit of how fast the function changes at x within the sphere of radius ϵ. Note that the function f itself cannot have a discontinuity, but the higher order derivatives may.

For a locally Lipschitz-continuous function the classical directional derivative need not exist. However, we can generalize the concept of differentiability by defining a generalized directional derivative as follows.

Definition 4.2 *Let $J : \mathbb{R}^n \to \mathbb{R}$ be a locally Lipschitz-continuous function at a point $x \in \mathbb{R}^n$. The generalized directional derivative of J at x in the direction $p \in \mathbb{R}^n$ is defined by*

$$J^{\mathrm{o}}(x; p) = \limsup_{y \to x \text{ as } t \downarrow 0} \frac{J(y + tp) - J(y)}{t},$$

where $y \in \mathbb{R}^n$ and $t \in \mathbb{R}$.

Note that the only difference between this definition and the definition of a traditional directional derivative is the sup, meaning that the largest directional derivative along any direction y is taken. At a differentiable point, these limits will be the same along any direction; however at a non-differentiable point these values may be different and choosing the largest is a choice of convenience that will be exploited later.

Unlike a traditional gradient, which is unique, at non-smooth points in general infinite subgradients exist as part of the *subdifferential* set, defined as follows.

Definition 4.3 *Let $J : \mathbb{R}^n \to \mathbb{R}$ be a locally Lipschitz-continuous function at a point $x \in \mathbb{R}^n$. Then the* subdifferential *of J at x is the set $\partial J(x)$ of vectors $\xi \in \mathbb{R}^n$ such that*

$$\partial J(\mathbf{x}) = \left\{ \xi \in \mathbb{R}^n | J^{\mathrm{o}}(x; p) \geq \xi^{\mathrm{T}} p \qquad \text{for all } p \in \mathbb{R}^n \right\}.$$

Each vector $\xi \in \partial J(x)$ is called a subgradient *of J at x.*

In analogy to the traditional interpretation of the gradient as a tangent hyperplane to the function J, intuitively one can consider a subgradient at x to be a normal vector of any tangent hyperplane that remains on or below the generalized directional derivatives $J^{\mathrm{o}}(x; p)$ in all directions p.

When J is differentiable at x there is only one element of the subdifferential $\partial J(x)$, and it is the standard gradient. Below, the gradients (12) and (16) will be replaced with a *subgradient* in order to allow for non-smooth optimization in general. Because the gradient and subgradient are identical at differentiable points, this change is primarily transparent to the model and adjoint development, although care must be taken to ensure that a value for the adjoint is chosen such that definition (4.3) holds at the discontinuities. In addition, care must be taken in the optimization algorithms. More details can be found in the book of Makela and Neittaanmaki (1992).

4.2. Optimization algorithms

Both 4D-Var and MLEF employ an optimization algorithm as an integral part of finding the solution. As shown in section 3, 4D-Var relies entirely upon this optimization algorithm while MLEF uses it to maximize the posterior probability density function at each step.

Following the derivation of Zupanski *et al.* (2008), we introduce two non-smooth iterative optimization methods.

In general, an iterative optimization algorithm can be formulated as

$$x_{k+1} = x_k + \alpha_k p_k, \tag{23}$$

where k is the iteration number, p_k is the search direction and α_k is the step length. This procedure continues until some convergence criterion has been met.

If $J(x_{k+1}) \leq J(x_k)$ for all k, then an iterative method is called a *descent method* and the direction p_k is called a *descent direction*. For smooth (continuously differentiable) objective functions, a descent direction may be generated by exploiting the fact that the direction opposite to the gradient is locally the direction of steepest descent. The step size α_k can then be determined, for example, using a line-search technique (see e.g. Nocedal and Wright, 2006). Furthermore, a necessary condition for local optimality is that the gradient goes to zero and by continuity becomes small on approach to an optimal point. This fact provides a useful stopping criterion for smooth iterative methods.

However, the direct application of smooth gradient-based methods to non-smooth problems may lead to a failure in convergence, optimality conditions or gradient approximation. The usage of subgradients allows us to generalize well-developed gradient-based methods for non-smooth problems. In this section, we detail the optimization algorithms used in this study. For the rest of the paper, a subgradient $\nabla J(x_k) \in \partial J(x_k)$ is denoted by ∇J_k.

4.3. L-BFGS

In this work, we test an implementation of the L-BFGS algorithm version VA15 of Liu and Nocedal (1989) in the Harwell library. The version of L-BFGS with subgradients is detailed in Zupanski *et al.* (2008).

The L-BFGS method is an adaptation of the BFGS method to large problems, achieved by changing the generalized Hessian update of the latter. The L-BFGS method uses an approximation H_k to an inverse generalized Hessian, which is updated at each time step. The search direction is found, in analogy to the Newton method, by

$$p_k = -H_k \nabla J_k. \tag{24}$$

An inverse generalized Hessian approximation is updated at each iteration by

$$H_{k+1} = V_k^{\mathrm{T}} H_k V_k + \rho_k s_k s_k^{\mathrm{T}}, \tag{25}$$

where $s_k = x_k - x_{k-1}$,

$$V_k = I - \rho_k y_k s_k^{\mathrm{T}}, \tag{26}$$

$y_k = \nabla J_k - \nabla J_{k-1}$ and $\rho_k = 1/(y_k^{\mathrm{T}} s_k)$.

In the L-BFGS method, instead of forming the matrices H_k explicitly (which would require a prohibitively large allocation of memory for even a medium-size problem), one only stores the vectors s_k and y_k obtained in the last m iterations, which define H_k implicitly; a cyclical procedure is

used to retain the latest vectors and discard the oldest ones. Thus, after the first m iterations, Eq. (25) becomes

$$\begin{aligned} H_{k+1} =& \Pi_{k-\hat{m}}^{\mathrm{T}} H_0 \Pi_{k-\hat{m}} \\ &+ \rho_{k-\hat{m}} \Pi_{k-\hat{m}+1}^{\mathrm{T}} S_{k-\hat{m}} \Pi_{k-\hat{m}+1} \\ &+ \rho_{k-\hat{m}+1} \Pi_{k-\hat{m}+2}^{\mathrm{T}} S_{k-\hat{m}+1} \Pi_{k-\hat{m}+2} \\ &\vdots \\ &+ \rho_k S_k, \end{aligned} \tag{27}$$

where $\Pi_j = V_j V_{j+1} \cdots V_k$, $S_i = s_i s_i^{\mathrm{T}}$, $\hat{m} = \min(k, m-1)$ and the initial approximation H_0 is taken to be I.

Thus, only at most $2m$ correction pairs s_i and y_i for $i = 1, \ldots, m$ are needed and no full matrix is ever stored in memory as Eq. (24) is solved.

4.4. LMBM

The L-BFGS method was originally created for smooth optimization and is generalized for non-smooth methods by replacing the gradient with an arbitrary subgradient. However, there are some potentially serious drawbacks to this approach. Firstly, it is theoretically possible for a non-descent search direction to occur, as the direction opposite to an arbitrary subgradient does not guarantee descent. Thus, a smooth line-search algorithm, used for step-size (α_k) selection, may fail. Secondly, due to the fact that the norm of an arbitrary subgradient does not necessarily become small in the neighbourhood of an optimal point, a convergence criterion based on this assumption, valid for smooth gradients, will also fail when the optimal point occurs at a discontinuity. Moreover, in general the convergence speed of subgradient methods can be poor.

In this subsection we describe the LMBM (Haarala *et al.*, 2004, 2007), where the above-mentioned drawbacks are avoided by using a so-called bundling technique (Lemarechal, 1975). The idea of bundling is that instead of using just one arbitrary subgradient we approximate the whole subdifferential (see Definition (4.3)) of the objective function by gathering the subgradients from previous iterations into a bundle. In this way, we obtain more information about the local behaviour of the function than an individual arbitrary subgradient can yield.

LMBM was specifically developed for solving large-scale non-smooth optimization problems. It is characterized by the usage of null steps together with a simple aggregation of subgradients. Moreover, as in L-BFGS, the limited-memory approach is utilized in the calculation of the search direction

$$p_k = -H_k \nabla \tilde{J}_k, \tag{28}$$

where $\nabla \tilde{J}_k \in \mathbb{R}^n$ is an aggregate subgradient and H_k is not formed explicitly but calculated by the L-BFGS update (see Eq. (27)) after a serious step and by the L-SR1 update (see e.g. Byrd *et al.* (1994)) after a null step. The usage of null steps gives further information about the non-smooth objective in the case in which the search direction is not 'good enough'. That is, a null step is taken when the descent criterion

$$J(x_k + t_R^k p_k) \leq J(x_k) + \varepsilon_R^k \tag{29}$$

is not satisfied. Here t_R^k is the step size and $\varepsilon_R^k > 0$ is the desired descent of J at x_k. In the case of a null step, we set $x_{k+1} = x_k$ but information about the objective function is increased by storing the auxiliary point $y_{k+1} = x_k + t_R^k p_k$ and the corresponding auxiliary subgradient $\nabla J_{k+1} \in \partial f(y_{k+1})$. These values are used in the computation of the new aggregate subgradient that is used in the next iteration. A simple aggregation of subgradients guarantees the global convergence of the method (for more details see Haarala *et al.*, 2007) and make it possible to evaluate a subgradient-based stopping criterion.

The pseudo-code of the LMBM algorithm is shown in Table II.

Table II. LMBM pseudo-code.

```
Program LMBM
  Initialize x_1 ∈ R^n, ∇J_1 ∈ ∂J(x_1), and ε_s > 0;
  Set k = 1 and p_1 = −∇J_1;
  While the termination condition
        w_k ≤ ε_s is not met
    Find step sizes t_L^k and t_R^k;
    Set x_{k+1} = x_k + t_L^k p_k
    Evaluate J(x_{k+1}) and ∇J_{k+1} ∈ ∂J(x_k + t_R^k p_k);
    If t_L^k > 0 then
      Compute the search direction p_{k+1}
        using ∇J_{k+1} and the L-BFGS update;
    Else
      Compute the aggregate
        subgradient ∇J̃_{k+1};
      Compute the search direction p_{k+1}
        using ∇J̃_{k+1} and the L-SR1 update;
    End if
    Set k = k + 1;
  End While
  Return final solution x_k;
End LMBM
```

5. Experimental set-up

We now consider data assimilation of the shallow-water equations with a discontinuous observation operator detailed below.

Starting from the exact initial conditions given in (5) and boundary conditions given in (2), the model is evolved forward in time. Observations of the model state (u, v, ϕ) are taken at each time step and every spatial grid point using an observation operator and are then perturbed with uncorrelated Gaussian noise of mean 0 and standard deviation $\sigma_{u_{\mathrm{obs}}}$, $\sigma_{v_{\mathrm{obs}}}$ and $\sigma_{\phi_{\mathrm{obs}}}$, respectively. The exact initial conditions are perturbed with correlated Gaussian noise of mean 0 and covariance matrix B. The perturbed initial conditions are the *background* and the problem is to reconstruct the exact solution optimally using background and observations.

Both the 4D-Var and MLEF approaches are used to solve this problem.

In this research, observations at all grid points are available for each time step –perhaps the best possible scenario for data assimilation. This removes the issue of sparsity of observations from the experimental set-up in order to focus on the impact of non-smooth observation operators. Thus,

assuming there are N, M non-boundary grid points in the x and y direction, respectively, $N_{\text{state}} = N_{\text{obs}} = 3NM$.

5.1. *Observation-error covariance matrix*

For this experiment, the observation-error covariance matrix R is taken to be diagonal. Thus, R^{-1} is the diagonal inverse matrix with

$$R_{i,i}^{-1} = \begin{cases} 1/\sigma_{u_{\text{obs}}}^2, & 1 \le i \le MN, \\ 1/\sigma_{v_{\text{obs}}}^2, & MN+1 \le i \le 2MN, \\ 1/\sigma_{\phi_{\text{obs}}}^2, & 2MN+1 \le i \le 3MN. \end{cases} \tag{30}$$

In this experiment, $\sigma_{u_{\text{obs}}} = \sigma_{v_{\text{obs}}} = 1\,\text{m}\,\text{s}^{-1}$ and $\sigma_{\phi_{\text{obs}}} = 12\,\text{m}^2\,\text{s}^2$ were chosen based on approximate geostrophic considerations. A sample of this uncorrelated noise at time t_0 is shown in Figure 3.

5.2. *Observation operator*

In this section we detail an observation operator with varying levels of non-smoothness in its components. This operator is not based on physical considerations but rather chosen solely to demonstrate the behaviour of the optimization algorithms.

The observation operator $\mathcal{H}$ is given by

$$\mathcal{H}(x_i) = \begin{cases} \mathcal{H}_1(u_i), & 1 \le i \le MN, \\ \mathcal{H}_2(v_{i-MN}), & MN+1 \le i \le 2MN, \\ \mathcal{H}_3(\phi_{i-2MN}), & 2MN+1 \le i \le 3MN, \end{cases} \tag{31}$$

$$\mathcal{H}_1(u_i) = \begin{cases} u_i^3/u_{\min}^2, & u_i < u_{\min}, \\ u_i^2/u_{\max}, & u_i \ge u_{\max}, \\ u_i, & \text{else}, \end{cases} \tag{32}$$

$$\mathcal{H}_2(v_i) = \begin{cases} \log(v_i + \delta), & v_i \ge 0, \\ \log(-v_i + \delta), & v_i < 0, \end{cases} \tag{33}$$

$$\mathcal{H}_3(\phi_i) = \begin{cases} \phi_i, & \phi_i < H_{\max}, \\ \phi_i^2/H_{\max}, & \phi_i \ge H_{\max}. \end{cases} \tag{34}$$

Here, $u_{\min} = -5\,\text{m}\,\text{s}^{-1}$, $u_{\max} = 5\,\text{m}\,\text{s}^{-1}$ and $H_{\max} = 20\,000$ m.

These components of the observation operator are shown in Figure 4 and the observation operator of the initial state plus observational noise is shown in Figure 5.

Due to the kinks in the observation operators, this becomes a non-smooth optimization problem. The discontinuity in the piecewise derivatives of $\mathcal{H}_1$–$\mathcal{H}_3$ are shown in Figure 6. Note that the discontinuity becomes progressively more acute in $\nabla\mathcal{H}_3$, $\nabla\mathcal{H}_1$ and $\nabla\mathcal{H}_2$, giving flexibility for testing the behaviour of non-smooth optimization algorithms. In $\nabla\mathcal{H}_2$, the parameter δ controls the size of the discontinuity. All of these functions are locally Lipschitz-continuous, and the best global Lipschitz constant for $\mathcal{H}_2$ is $1/\delta$.

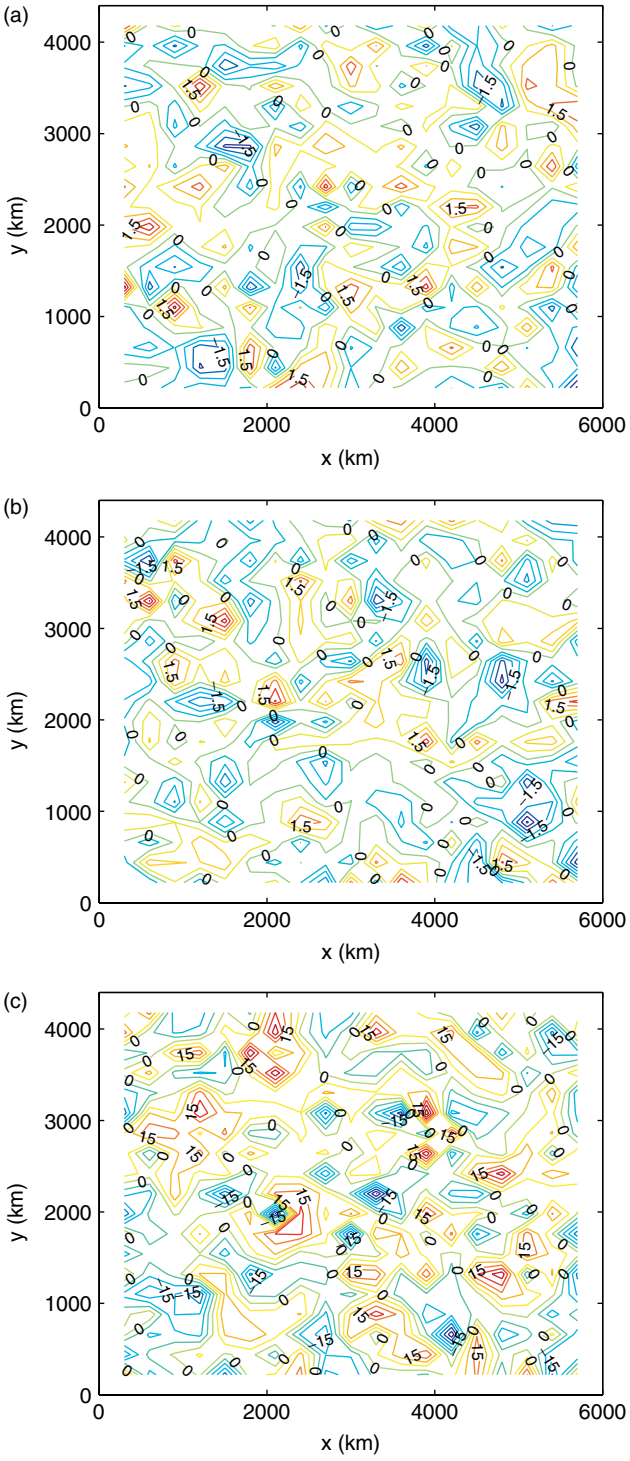

Figure 3. Observation error sample at t_0 for u, v and ϕ. (a) $\eta_{\text{obs}}^{(u)}$ sample, contour interval 0.5; (b) $\eta_{\text{obs}}^{(v)}$ sample, contour interval 0.5; (c) $\eta_{\text{obs}}^{(\phi)}$ sample, contour interval 5. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

5.3. *Background-error covariance matrix*

In this work we use an exact background-error covariance matrix to perturb the background x_{b}. The perturbation to the background vector is given by

$$\begin{aligned} u_{\text{b}} &= u_0 + \eta_u, \\ v_{\text{b}} &= v_0 + \eta_v, \\ \phi_{\text{b}} &= \phi_0 + \eta_\phi, \end{aligned} \tag{35}$$

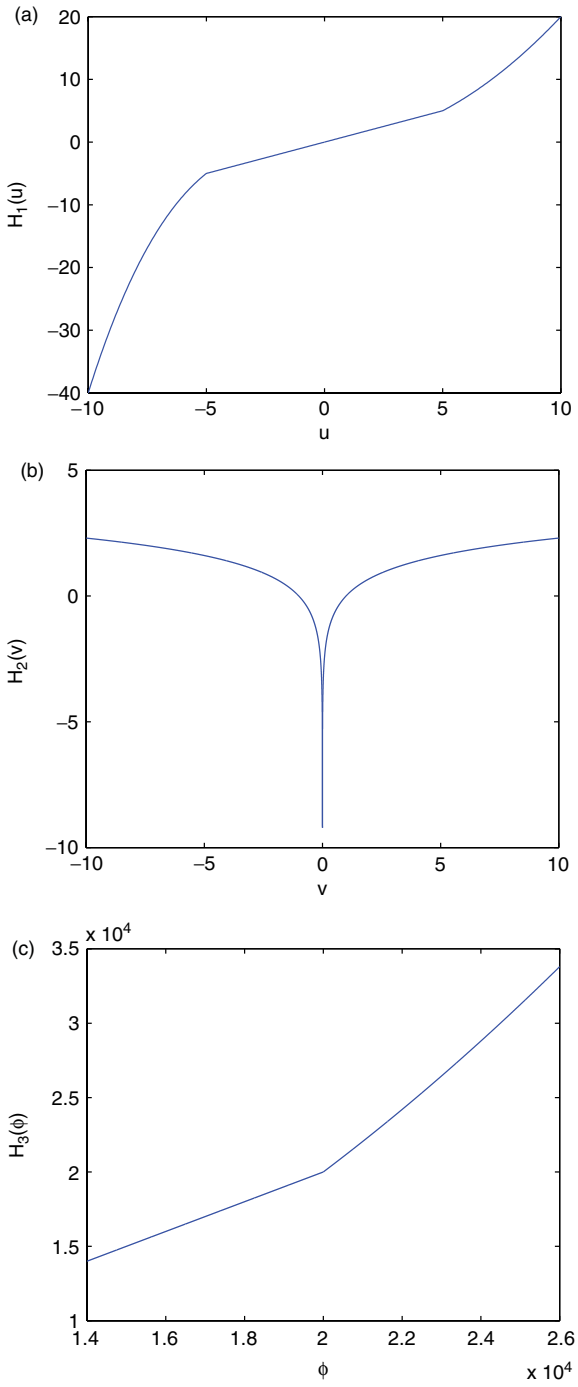

Figure 4. Observation operators: (a) $\mathcal{H}_1(u)$, (b) $\mathcal{H}_2(v)$ and (c) $\mathcal{H}_3(\phi)$. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

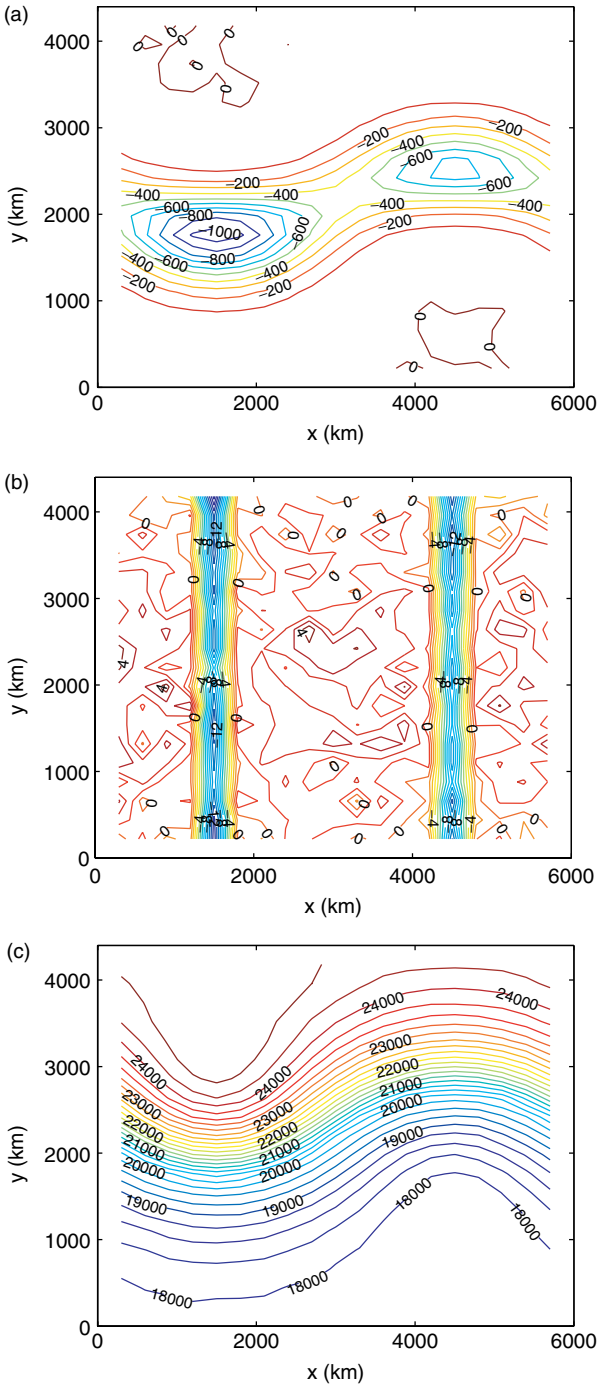

Figure 5. Observation sample for u, v, and ϕ at t_0. (a) u_{obs}, contour interval 100. (b) v_{obs}, contour interval 1; the two columns of tightly spaced contours are caused by the sharp drop in $\mathcal{H}_2(v)$ at $v = 0$. (c) ϕ_{obs}, contour interval 250. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

where (u_i, v_i, ϕ_i) are the exact initial conditions given in (5), $\eta_u \sim N(0, \Sigma_u)$, $\eta_v \sim N(0, \Sigma_v)$ and $\eta_\phi \sim N(0, \Sigma_\phi)$, where Σ_u, Σ_v, $\Sigma_\phi \in \mathbb{R}^{NM\times NM}$ are the covariance matrices of (u, v, ϕ). Thus, B is the block matrix

$$B = \begin{pmatrix} \Sigma_u & 0 & 0 \\ 0 & \Sigma_v & 0 \\ 0 & 0 & \Sigma_\phi \end{pmatrix} \tag{36}$$

and

$$B^{-1} = \begin{pmatrix} \Sigma_u^{-1} & 0 & 0 \\ 0 & \Sigma_v^{-1} & 0 \\ 0 & 0 & \Sigma_\phi^{-1} \end{pmatrix}. \tag{37}$$

$\Sigma_u = \sigma_u^2\Sigma$, $\Sigma_v = \sigma_v^2\Sigma$ and $\Sigma_\phi = \sigma_\phi^2\Sigma$ are created by the exponential squared kernel of

$$\Sigma = B^{1/2}B^{\mathrm{T}/2} \tag{38}$$

and

$$B^{1/2}_{i,j} = B^{1/2}_{j,i} = \exp(-r^2_{i,j}/L^2). \tag{39}$$

Here, r represents the Euclidean distance between two points in the grid, i.e.

$$r^2_{i,j} = \Delta x^2_{i,j} + \Delta y^2_{i,j}, \tag{40}$$

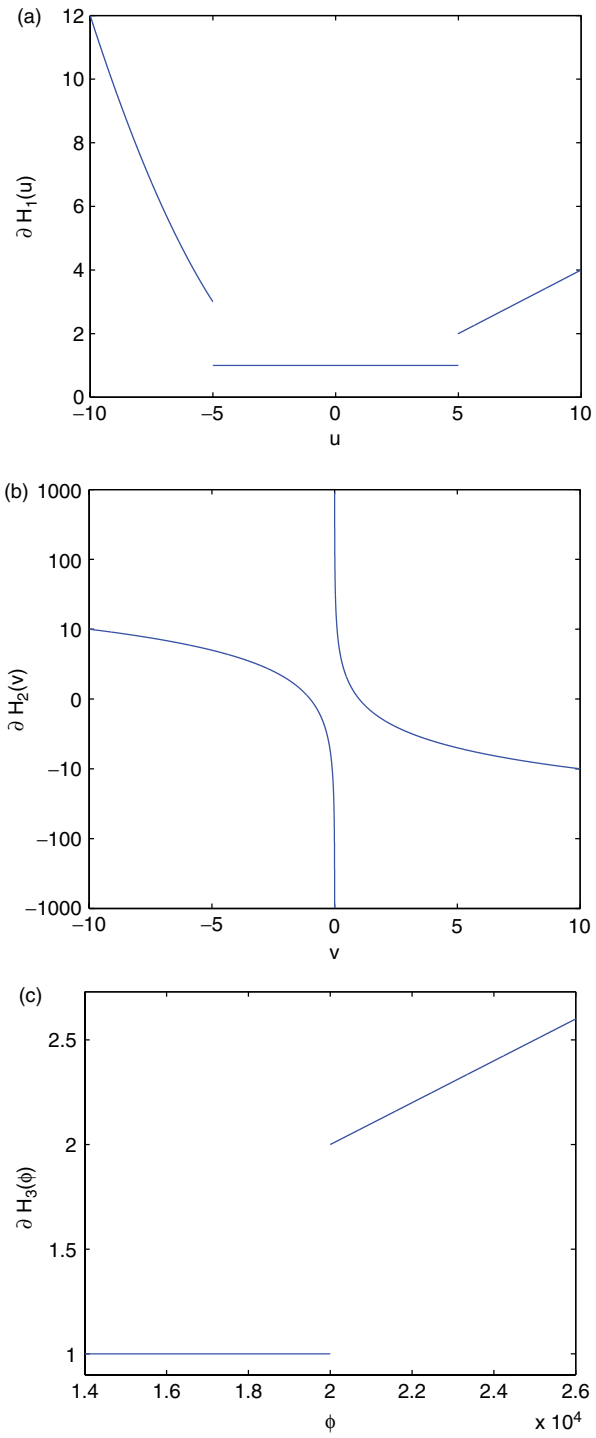

Figure 6. Piecewise derivative of observation operator: (a) $\nabla\mathcal{H}_1(u)$; (b) $\nabla\mathcal{H}_2(v), \delta = 10^{-3}$; (c) $\nabla\mathcal{H}_3(\phi)$. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

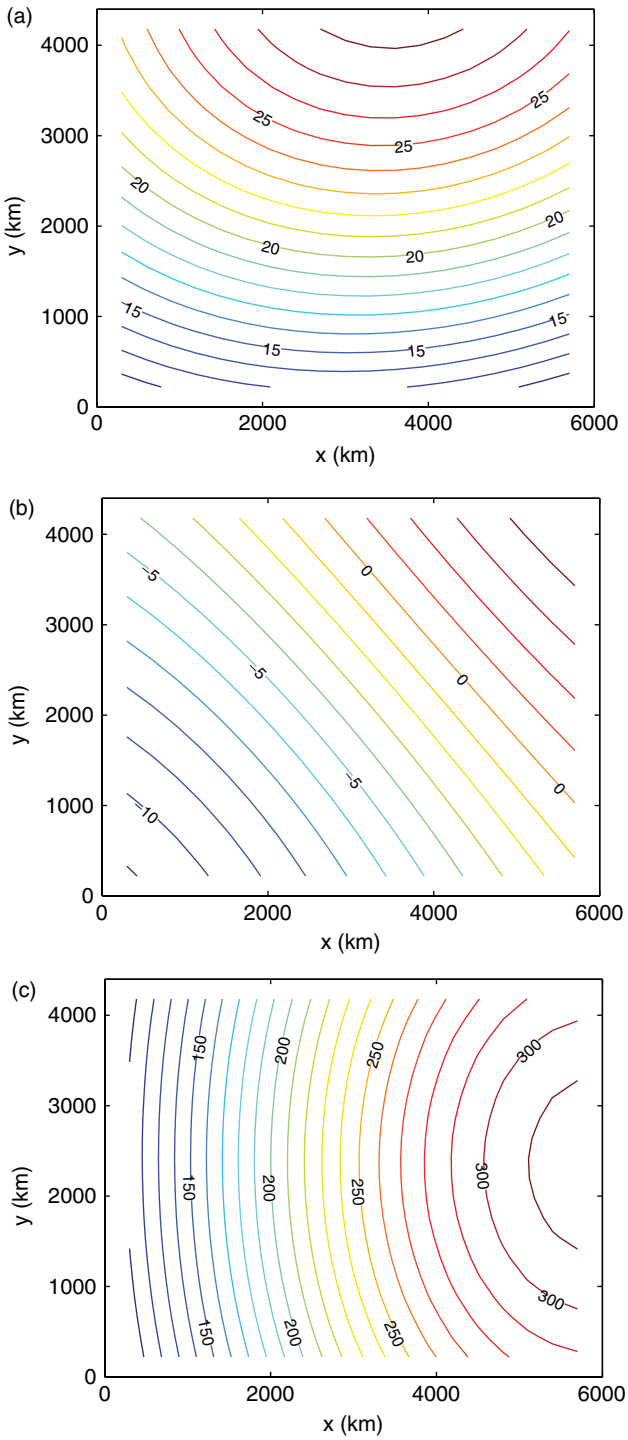

Figure 7. Background perturbation versus exact solution: (a) u_b, contour interval 1; (b) v_b, contour interval 1; (c) ϕ_b, contour interval 10. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

where $\Delta x_{i,j}$ and $\Delta y_{i,j}$ are the x- and y-coordinate distances in metres between the two points with global indexes i and j, respectively. L represents the correlation length in metres required for the correlation between two points to reach $1/\mathrm{e} \approx 0.36788$.

The normal perturbation η_u is created with the transformation $\eta_u = \Sigma_u^{1/2} Z_u$, where $Z_u \in \mathbb{R}^{NM}$ and the components of $Z_{u_i} \sim N(0,1)$ are independent and identically distributed standard normal variables. η_v and η_ϕ are created the same manner with Σ_v, Σ_ϕ, Z_v and Z_ϕ uncorrelated standard normal variables, respectively.

Based on the results in Zupanski *et al.* (2006), a correlation length of $L = 7000$ km and background-perturbation magnitudes of $\sigma_u = \sigma_v = 20$ m and $\sigma_\phi = 200\,\mathrm{m^2\,s^{-2}}$ were chosen. As discussed below, the background used in the 4D-Var experiment and MLEF experiments is found by taking an ensemble average over 96 realizations of (u_b, v_b, ϕ_b). The actual perturbation versus the exact solution used as the background is shown in Figure 7.

While the matrix B has, in theory, full rank, like many other background-error covariances used in practice, this background-error covariance matrix is numerically

rank-deficient due to the effect of having a much larger correlation length than the grid spacing. Thus, the inverse background-error covariance cannot be computed numerically. This issue is addressed in our 4D-Var implementation with the approach used in Chen *et al.* (2011), i.e. by applying the change-of-variables transformation $\delta_{\mathrm{b}}(x) = x - x_{\mathrm{b}} = B^{1/2}z$ and using z as the control variable. Since $B^{1/2}B^{\mathrm{T}/2} = B$, this removes the necessity of obtaining the inverse from the cost function, so that (8) becomes

$$J(z) = \frac{1}{2}z^{\mathrm{T}}z + \frac{1}{2}\sum_{k=0}^{NT}\delta'_{y_k}(z)^{\mathrm{T}}R^{-1}\delta'_{y_k}(z), \tag{41}$$

where $\delta'_{y_k}(z') = \delta_{y_k}(B^{1/2}z + x_{\mathrm{b}})$, while the gradient (Chen *et al.*, 2011) becomes

$$\nabla'_z J(z') = z' - \sum_{k=0}^{NT}\frac{\partial x^{(k)}}{\partial z'}\frac{\partial \mathcal{H}}{\partial x^{(k)}}R^{-1}\delta'_{y_k}, \tag{42}$$

and, as $x(0) = B^{1/2}z + x_{\mathrm{b}}$,

$$\frac{\partial x^{(0)}}{\partial z}(z) = B^{1/2}. \tag{43}$$

5.4. Design

In order to compare both 4D-Var and MLEF in the presence of non-smooth observation operators, the following procedure is used.

(1) The initial condition for h is listed in Eq. (4). From this, u_0, v_0 and ϕ_0 are created by (5).
(2) These conditions are evolved forward in time by solving (1) through (3) to create $x_{\mathrm{exact}} = (u_{\mathrm{exact}}, v_{\mathrm{exact}}, \phi_{\mathrm{exact}})$ at times $t = 0, \ldots, NT$.
(3) Observations: Gaussian noise, as detailed in section 5.1, is added to $\mathcal{H}(x_{\mathrm{exact}})$ to create the observations.

MLEF:

(4a) Ensemble members: each member of the initial ensemble is created by sampling from a correlated Gaussian random variable with mean x_{exact} and covariance matrix B, as discussed in section 5.3.
(5a) Control state: the mean of the ensemble created in (4a) is used as the control state.

4D-Var:

(4b) Background value: the MLEF initial control state from (5a) is used as the background value.
(5b) Background-error covariance matrix: the matrix B used in (4a) to create the ensemble is used. The matrix $B^{1/2}$ is found using an eigenvalue decomposition.

Both MLEF and 4D-Var:

(6) The experiment is then run and the root-mean-squared error (RMSE) is taken at each time step versus $(u_{\mathrm{exact}}, v_{\mathrm{exact}}, \phi_{\mathrm{exact}})$, as detailed in section 5.6.

Table III. Experimental set-up.

Experiment #	u obs op	v obs op	ϕ obsop
1	linear	linear	linear
2	linear	linear	H_3
3	H_1	linear	H_3
4	H_1	H_2	H_3

5.5. Experiments

Within the context above, we now design three numerical experiments to test the performance of L-BFGS and LMBM within the 4D-Var and MLEF frameworks in order to assess their performance in data assimilation in the presence of non-differentiable observation operators of increasing difficulty. The experimental set-up is shown in Table III. Experiment 1 is the most favuorable experimental set-up for data assimilation with all linear (and thus differentiable) observation operators. Experiment 2 possesses only a 'slight' non-smoothness in the observation operator for ϕ and experiment 3 has a discontinuity in both u and ϕ, while experiment 4 constitutes the most difficult case with a sharp discontinuity in the observation operator for v and the same operator as experiment 3 for u and ϕ.

5.6. Success criteria

We now define our success criteria for the experiments listed in section 5.5.

To judge the quality of the assimilation results, we use the RMSE of the calculated solution versus the exact solution. The RMSE for cycle (k) is calculated as follows:

$$\begin{aligned}
\mathrm{RMSE}_u^{(k)} &= \sqrt{\frac{\left(u_{\mathrm{exact}}^{(k)} - u^{(k)}\right)^{\mathrm{T}}\left(u_{\mathrm{exact}}^{(k)} - u^{(k)}\right)}{NM}},\\
\mathrm{RMSE}_v^{(k)} &= \sqrt{\frac{\left(v_{\mathrm{exact}}^{(k)} - v^{(k)}\right)^{\mathrm{T}}\left(v_{\mathrm{exact}}^{(k)} - v^{(k)}\right)}{NM}},\\
\mathrm{RMSE}_\phi^{(k)} &= \sqrt{\frac{\left(\phi_{\mathrm{exact}}^{(k)} - \phi^{(k)}\right)^{\mathrm{T}}\left(\phi_{\mathrm{exact}}^{(k)} - \phi^{(k)}\right)}{NM}}.
\end{aligned} \tag{44}$$

As is common in data-assimilation experiments, success is judged by the assimilation achieving an RMSE that is lower than both the observation and background errors. In this case, the expected RMSE from simply using the observations is much lower than that the expected RMSE of the background. Thus, in order for the data-assimilation procedure to be considered a success, the RMSE must reach a level lower than the observational noise $(\sigma_{u_{\mathrm{obs}}}, \sigma_{v_{\mathrm{obs}}}, \sigma_{\phi_{\mathrm{obs}}})$ from section 5.1.

6. Numerical results

In this section we present the numerical results for both MLEF and 4D-Var with the experimental set-up detailed in section 5.

6.1. *MLEF results*

For MLEF, 96 ensemble members are used with the experimental set-up detailed in section 5.4 to test experiments 1–4 using MLEF with the L-BFGS quasi-Newton method as well as the LMBM algorithm. Up to 100 iterations of the minimization algorithm are allowed at each step for both methods, with the minimization terminating at each step if the change between subsequent normalized cost-function evaluations is less than 10^{-1}. Note that in a large-scale implementation this would mean a prohibitively large number of observation-operator evaluations, but it was chosen here to focus on the actual performance of the minimization algorithms. The δ parameter, used in experiment 4 in the observation operator $\mathcal{H}_2(v)$, is set to $\delta = 10^{-4}$.

The RMSEs for the four experiments using MLEF and the smooth optimization algorithm are shown for the u and v components of the velocity field along with the geopotential field ϕ in Figure 8.

The results show a drop from the background state below the observation RMSE in just one cycle for experiments 1–3; however, experiment 4 is not successful, as the RMSE in v and ϕ never reach below the expected observational RMSE at the same time. In this experiment L-BFGS initially fails to reduce the RMSE of ϕ sufficiently much beyond the level of the background, while later in the assimilation process the RMSE in v grows above the level of the observation error. Both of these symptoms are indicative of problems in the minimization process owing to the non-smoothness of the observation operator. Note also that the RMSE oscillates between successful time steps. The reasons for this are twofold: firstly, unlike 4D-Var, MLEF does not take the model into consideration during the minimization process. Secondly, the numerical scheme for the shallow-water equation model uses a two time-step filtering scheme (Grammeltvedt, 1969). These two issues combined cause the MLEF solution to oscillate between two different trajectories based on the model fluctuations across time steps.

The results from experiments 2–3 demonstrate that, as predicted by theory in Zupanski (2005), MLEF can handle slightly non-smooth cases even with an algorithm originally designed for smooth optimization in place. However, the RMSE from experiment 4 shows that MLEF with the smooth L-BFGS algorithm has difficulty with a highly non-smooth data-assimilation case.

These experiments are repeated with the LMBM algorithm in place, and the results are shown in Figure 9. The results of using LMBM show that all experiments are now successful, even the most difficult cases. However, once again the RMSE level of ϕ reaches the expected RMSE from observations alone after approximately 50 cycles.

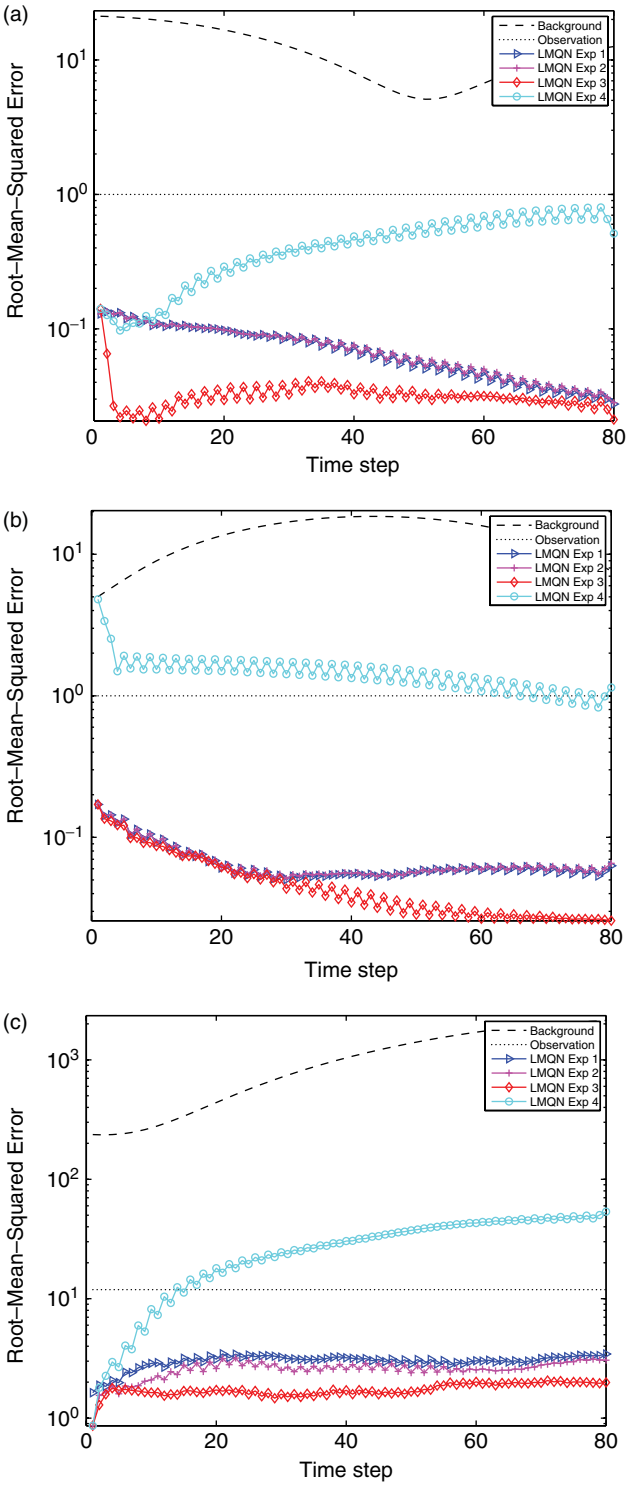

Figure 8. MLEF RMSE versus data-assimilation cycle, smooth minimization algorithm. Since observations are taken at each model time step, the cycle number is equal to the model time step. Thus the y-axis represents the RMSE achieved during each cycle. (a) u RMSE–L-BFGS; (b) v RMSE–L-BFGS; (c) ϕ RMSE–L-BFGS. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

6.2. *4D-Var results*

For 4D-Var, the four experiments from section 5.5 are run for the cost function (41) with the set-up described in section 5.4. As for the MLEF case, the δ parameter for $\mathcal{H}_2$ is set to $\delta = 10^{-4}$.

The performance of the L-BFGS and LMBM methods (measured in terms of cost-function value versus number of cost-function evaluations) is shown in Figure 10(a) and (b). The results demonstrate that the more challenging non-smooth experiments require more iterations. In addition, we see that L-BFGS fails to converge for experiment 4.

The results of using L-BFGS are shown in Figure 11, while the LMBM results are shown in Figure 12. For the first three experiments, the level of RMSE achieved for both methods is nearly identical and the cost difference between

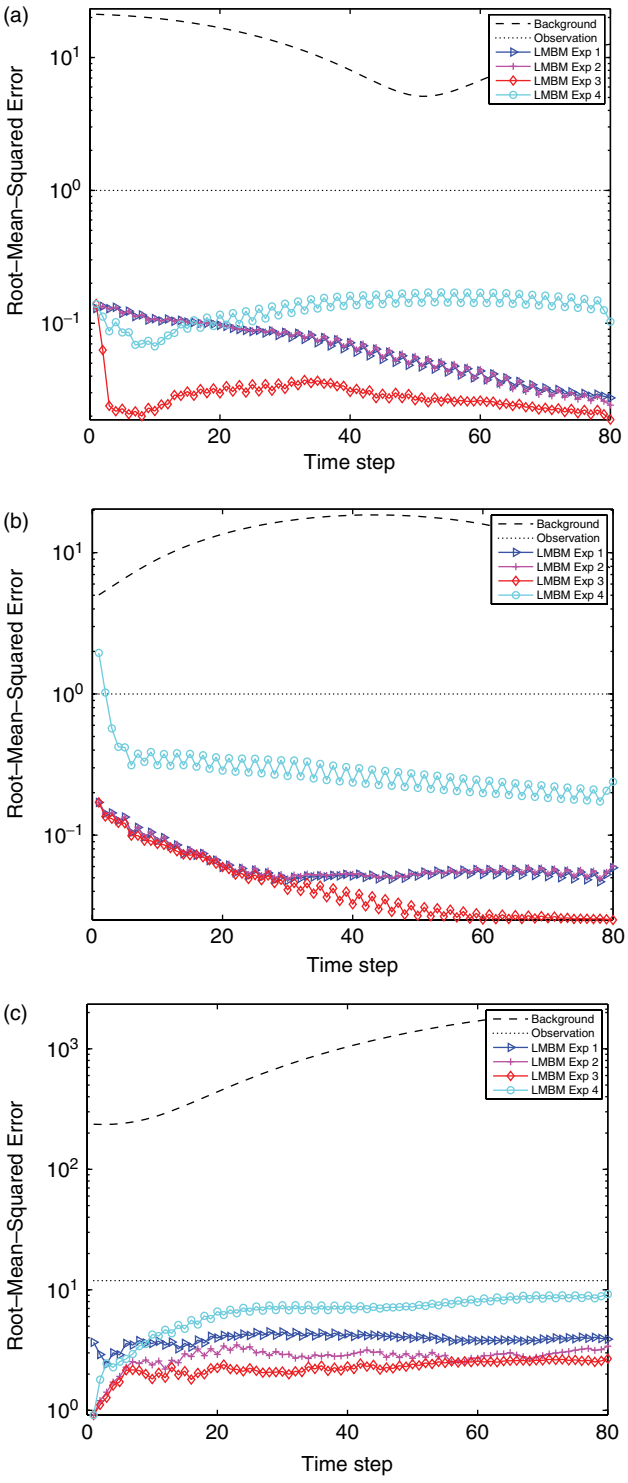

Figure 9. MLEF RMSE versus data-assimilation cycle, LMBM algorithm. (a) u RMSE–LMBM; (b) v RMSE–LMBM; (c) ϕ RMSE–LMBM. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

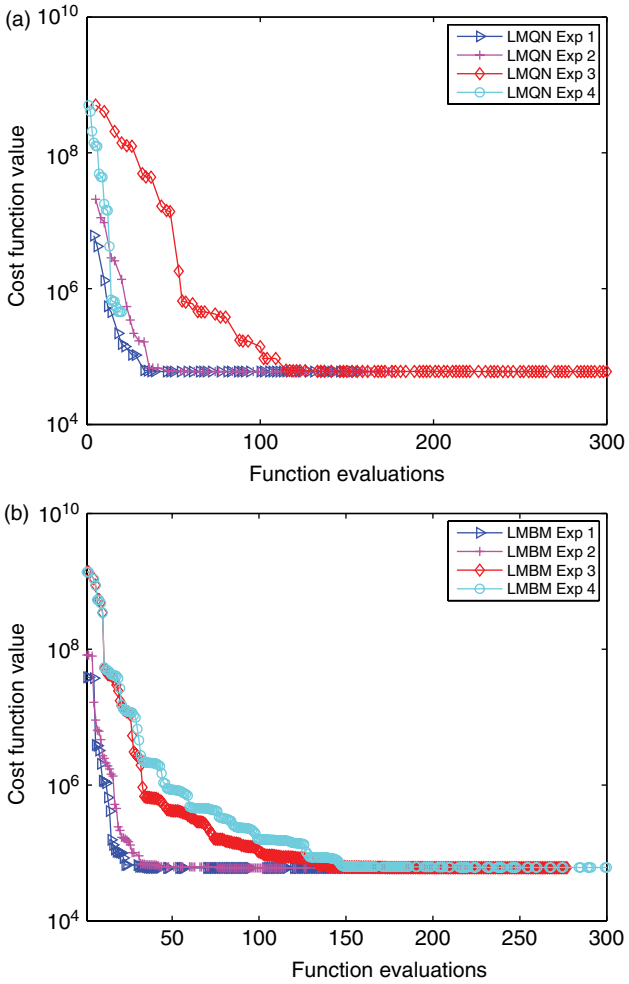

Figure 10. 4D-Var cost history. The x-axis shows the number of cost-function evaluations required for the minimization, while the y-axis shows the cost function achieved at that point in the optimization. (a) 4D-Var cost, L-BFGS; (b) 4D-Var cost, LMBM. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

L-BFGS and LMBM is not pronounced. However, for the fourth experiment L-BFGS does not successfully reduce the RMSE below that expected only from observations, and thus has failed on this challenging non-smooth case. Changing the line search from strong to weak Wolfe conditions, as suggested in Lewis and Overton (2008b), does not remedy the situation. LMBM, however, is able to handle this situation with the same level of accuracy as the other cases.

By adjusting the parameter δ, we can control the Lipschitz parameter of the observation operator $\mathcal{H}_2$, thus increasing the difficulty of the non-smooth optimization. The results of varying δ for both L-BFGS and LMBM are shown in Figures 13 and 14, respectively. These results show that LMBM can successfully handle even the case in which $\delta = 10^{-8}$.

The error in the final solution versus the exact solution found by LMBM for experiment 4 is shown in Figure 15 for $\delta = 10^{-4}$. The error is small and has evolved in time away from the smooth background error in Figure 7.

7. Conclusions

In this research, we tested the impact of non-differentiable observation operators on the data assimilation of a limited-area shallow-water equation model. By simply replacing the gradient of the cost function with the subgradient, both 4D-Var and MLEF are able to assimilate the non-smooth observations to varying degrees of success with a smooth optimization algorithm, especially when the non-smoothness is not severe, as is the case for experiment 2. However, both methodologies encounter difficulties with the more sharply non-smooth experiments 3 and 4. This difficulty can be remedied in both MLEF and 4D-Var with

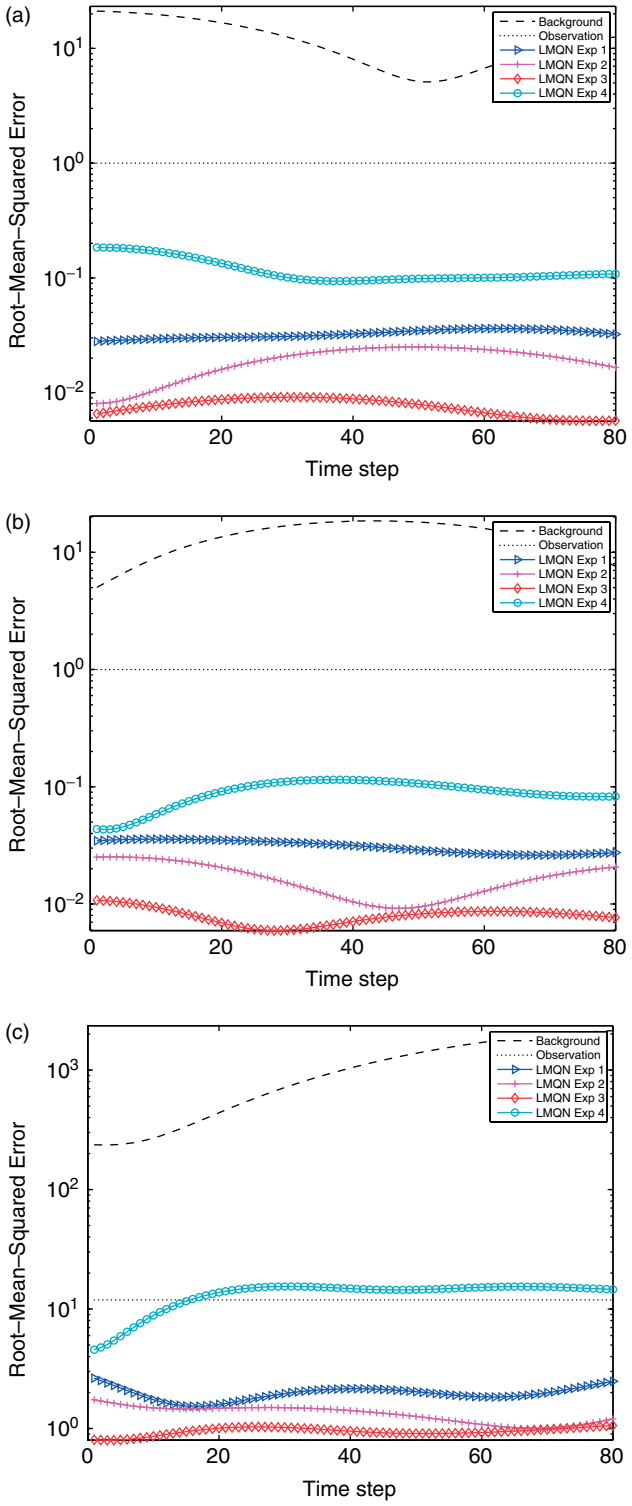

Figure 11. 4D-Var RMSE versus data-assimilation cycle, L-BFGS algorithm. (a) u RMSE; (b) v RMSE; (c) ϕ RMSE. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

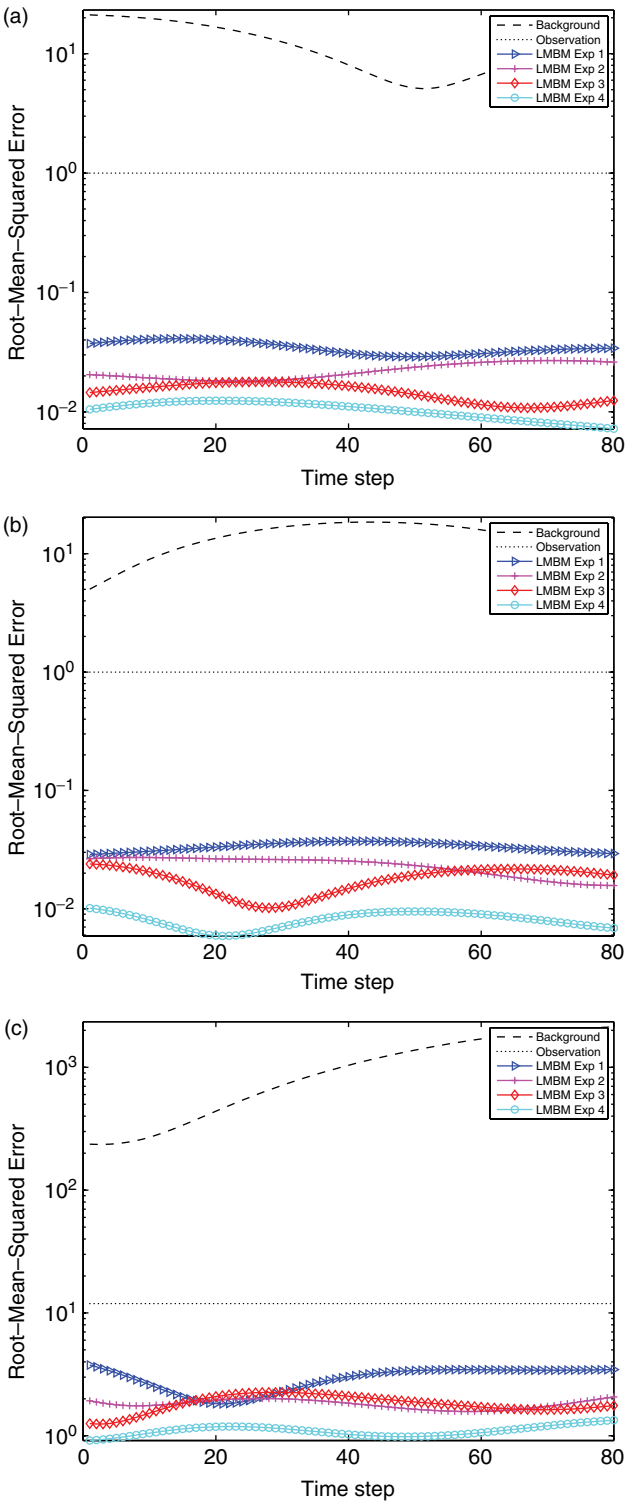

Figure 12. 4D-Var RMSE versus data-assimilation cycle, LMBM algorithm. (a) u RMSE; (b) v RMSE; (c) ϕ RMSE. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

the use of an algorithm specifically designed for non-smooth optimization, which in this research was the LMBM.

As MLEF is a sequential algorithm, it is not able to incorporate observations beyond a single time step. Like all data-assimilation algorithms that do not use the model as a constraint, this can lead to solutions that satisfy the minimization algorithm but are unphysical. 4D-Var, on the other hand, optimizes over the entire window, and thus can take advantage of observations at various times. In light of this, it is not surprising that 4D-Var achieves superior results in this particular case. However, one major strength of MLEF is that, unlike 4D-Var, the adjoint/tangent linear model of neither $\mathcal{M}$ nor $\mathcal{H}$ is required. For complex models, developing, testing and maintaining these adjoints represent a major investment. Since in practice data-assimilation cycles may not cover 80 cycles, as tested here, the results of MLEF, approximately equal to those of 4D-Var, may be more than sufficient.

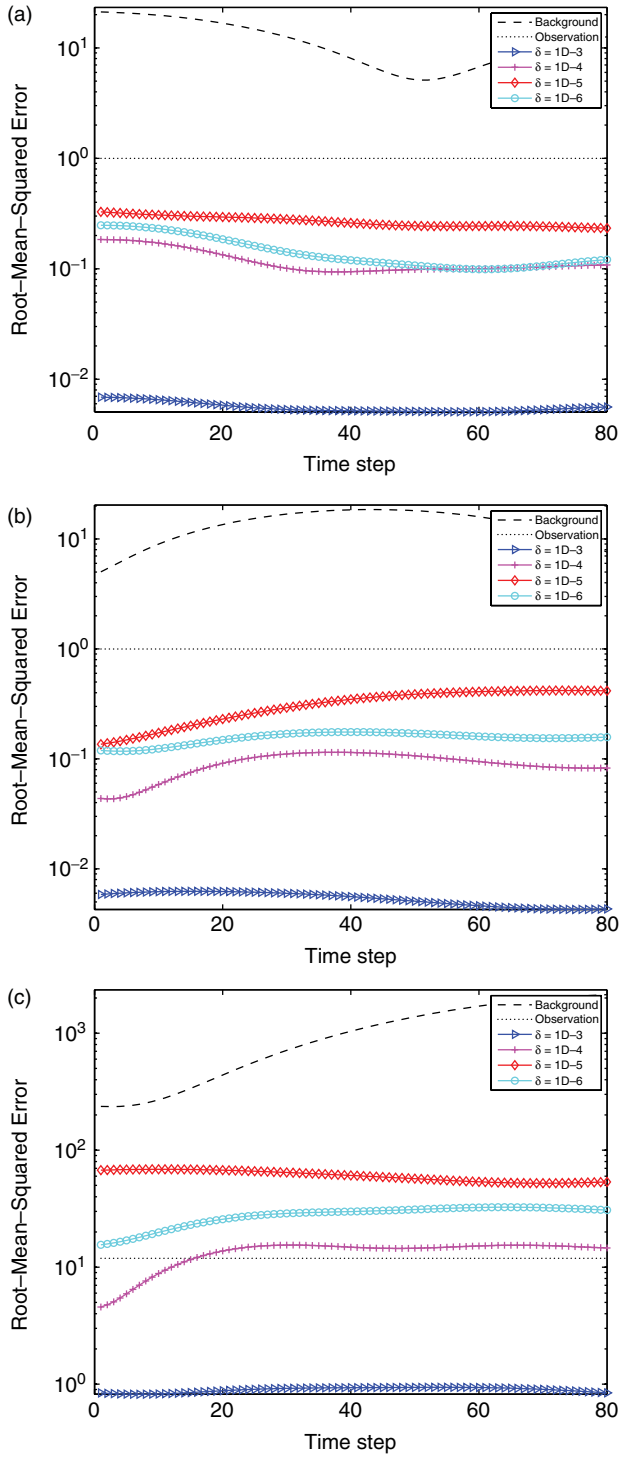

Figure 13. 4D-Var: impact of δ on RMSE versus data-assimilation cycle for experiment 4, L-BFGS. (a) u RMSE; (b) v RMSE; (c) ϕ RMSE. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

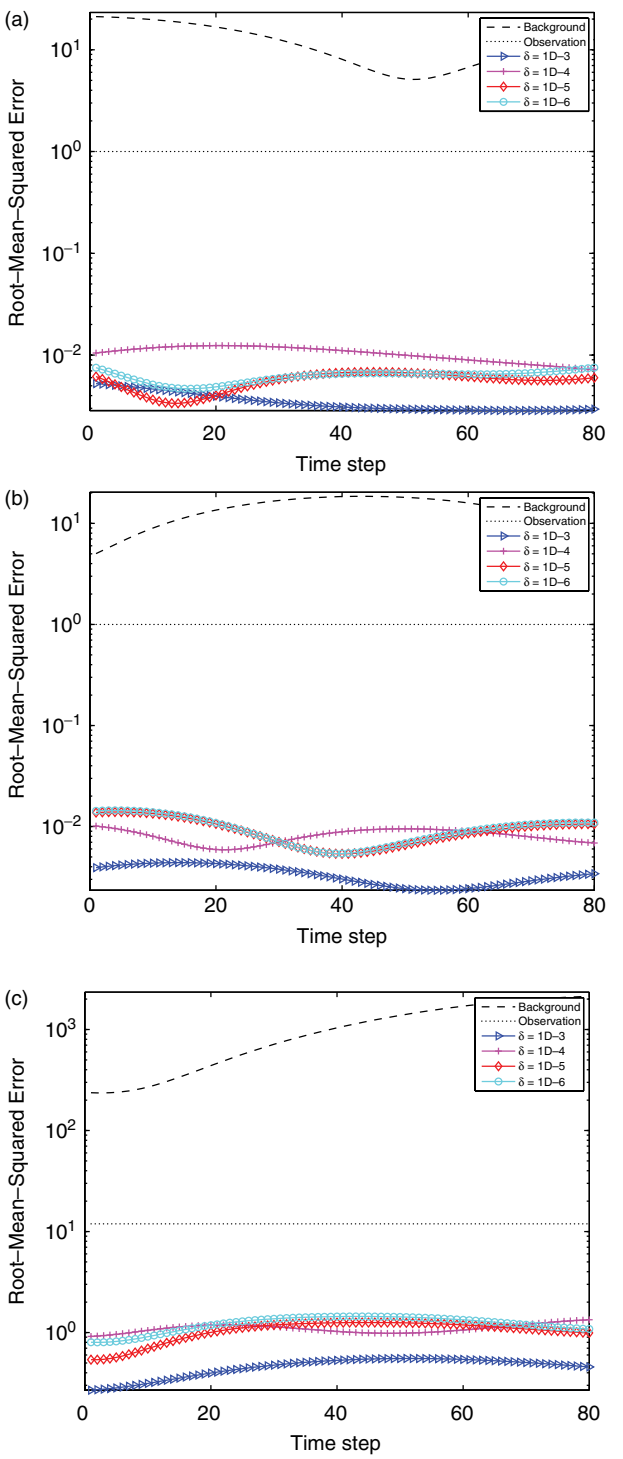

Figure 14. 4D-Var: impact of δ on RMSE versus data-assimilation cycle for experiment 4, LMBM. (a) u RMSE; (b) v RMSE; (c) ϕ RMSE. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

While the limited-memory bundle algorithm worked well in these cases, it should be noted that additional work remains to be done on LMBM to improve its numerical stability. With low error tolerances or numerically unstable cost functions and gradients, LMBM can often fail to achieve the desired minimization results (Karmitsa, 2007). In fact, in order to run the above minimization effectively, double-precision versions of MLEF and 4D-Var were required; single precision was not numerically stable enough for LMBM to perform successfully. This may be an issue for large-scale numerical weather prediction models that use single-precision variables.

L-BFGS handled the first three non-smooth cases well, but failed on the more difficult experiment 4 for $\delta < 10^{-3}$. This translates into an observation operator with a Lipschitz constant greater than 1000. This suggests that L-BFGS performs well when the Lipschitz constants are not extreme. In particular, LMBM, which pays careful attention to line-search and convergence issues, enables it to perform successfully in practice far beyond the range at which

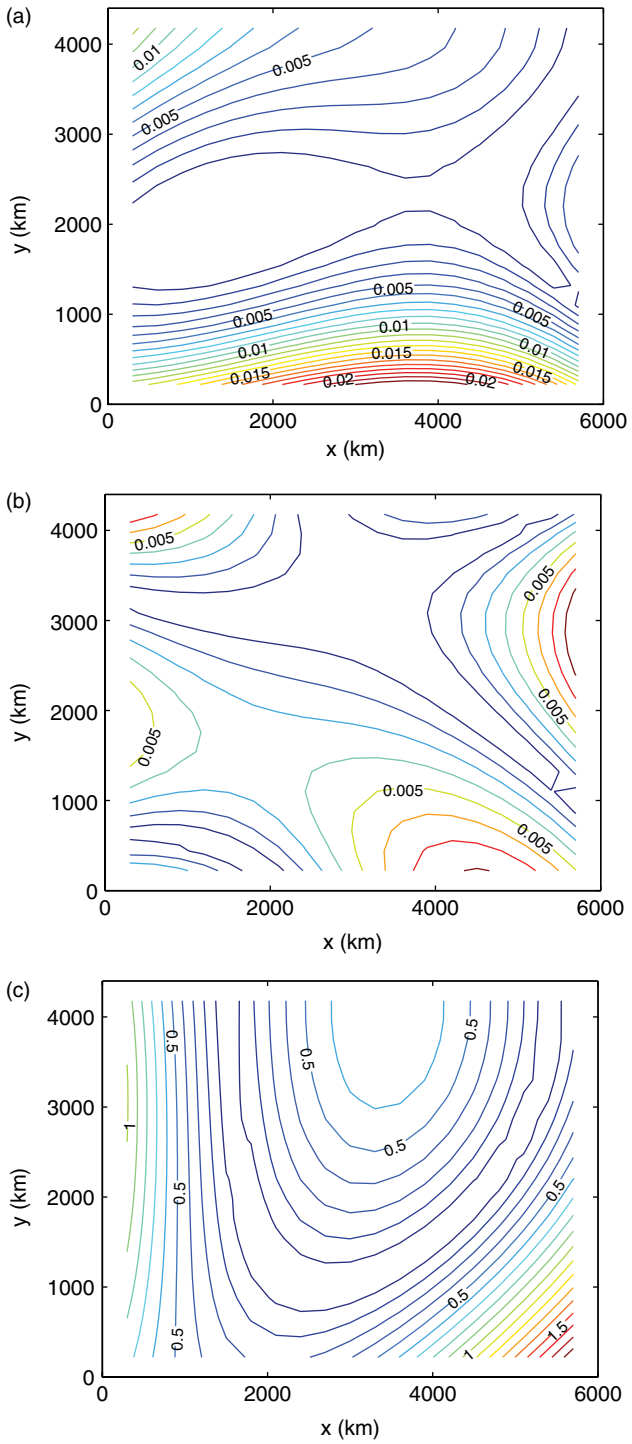

Figure 15. Final error of the computed versus exact solution. (a) u, contour interval 10^{-3}; (b) v, contour interval 10^{-3}; (c) ϕ, contour interval 10^{-1}. This figure is available in colour online at wileyonlinelibrary.com/journal/qj

L-BFGS fails. The use of 'null steps', which do not progress the optimization algorithm but only contribute additional information about the function, allows LMBM to handle such difficult cases. In addition, the use of a modified line search and avoiding a convergence criterion based on small gradients enables LMBM to be globally convergent. While it is possible to have some measure of success without paying attention to these issues, as shown the adverse effects on data assimilation become increasingly apparent at larger Lipschitz constants. A globally convergent line search such as the null-step approach used by LMBM is thus recommended for non-smooth data assimilation with large Lipschitz constants.

While data assimilation of non-smooth observation operators using this model –with control variables of the order of 10^3 –was successful, it remains to be seen whether similar results may be obtained in the case of data assimilation using realistic non-smooth observation operators and an actual operational weather prediction model with number of variables of the order of 10^7. Continued research in this area is needed in order to be of practical benefit to operational weather prediction centres and other large-scale data-assimilation optimal control problems. In particular, the assimilation of all-sky satellite radiances, long considered the holy grail of atmospheric data assimilation, may benefit greatly from continued research in this direction.

Acknowledgements

Professor I. M. Navon and Jeff Steward acknowledge the support of NSF/CMG grant ATM-0931198. Professor Zupanski acknowledges the NSF/CMG grant ATM-0930265. Professor Karmitsa acknowledges the support of the University of Turku (Finland). We also thank two anonymous reviewers for their helpful comments and suggestions. Finally, we would like to acknowledge high-performance computing support provided by NCAR's Computational and Information Systems Laboratory, sponsored by the National Science Foundation.

References

Bardos C, Pironneau O. 2005. Data assimilation for conservation laws. *Methods and Applications of Analysis* **12**: 103.

Bishop CH, Etherton BJ, Majumdar SJ. 2001. Adaptive sampling with the ensemble transform Kalman filter. Part I: theoretical aspects. *Mon. Weather Rev.* **129**: 420–436.

Byrd RH, Nocedal J, Schnabel RB. 1994. Representations of quasi-Newton matrices and their use in limited memory methods. *Mathematical Programming* **63**: 129–156. DOI: 10.1007/BF01582063

Chen X, Akella S, Navon IM. 2011. A dual-weighted trust-region adaptive POD 4-D Var applied to a finite-volume shallow-water equations model on the sphere. *Int. J. Numer. Methods Fluids* DOI: 10.1002/fld.2523

Chevallier F, Lopez P, Tompkins A, Janiskova M, Moreau E. 2004. The capability of 4D-Var systems to assimilate cloud-affected satellite infrared radiances. *Q. J. R. Meteorol. Soc.* **130**: 917–932. DOI: 10.1256/qj.03.113

Errico RM, Bauer P, Mahfouf JF. 2007. Issues regarding the assimilation of cloud and precipitation data. *J. Atmos. Sci.* **64**: 3785–3798. DOI: 10.1175/2006JAS2044.1

Fletcher SJ, Zupanski M. 2008. A study of ensemble size and shallow water dynamics with the maximum likelihood ensemble filter. *Tellus A* **60**: 348–360. DOI: 10.1111/j.1600-0870.2007.00294.x

Grammeltvedt A. 1969. A survey of finite-difference schemes for the primitive equations for a barotropic fluid. *Mon. Weather Rev.* **97**: 384. DOI: 10.1175/1520-0493(1969)097⟨0384:ASOFSF⟩2.3.CO;2

Greenwald TJ, Vukicevic T, Grasso LD, Vonder Haar T. 2004. Adjoint sensitivity analysis of an observational operator for visible and infrared cloudy-sky radiance assimilation. *Q. J. R. Meteorol. Soc.* **130**: 685–705. DOI: 10.1256/qj.03.44

Gunzburger MD. 2003. *Perspectives in flow control and optimization*. Society for Industrial Mathematics: Philadelphia, PA; ISBN 089871527X.

Haarala M, Miettinen K, Makela MM. 2004. New limited memory bundle method for large-scale nonsmooth optimization. *Optimization Methods and Software* **19**: 673. DOI: 10.1080/10556780410001689225

Haarala N, Miettinen K, Makela MM. 2007. Globally convergent limited memory bundle method for large-scale nonsmooth optimization. *Mathematical Programming* **109**: 181–205. DOI: 10.1007/s10107-006-0728-2

Homescu C, Navon IM. 2003. Optimal control of flow with discontinuities. *J. Comput. Phys.* **187**: 660–682. DOI: 10.1016/S0021-9991(03)00154-2

Honda Y, Nishijima M, Koizumi K, Ohta Y, Tamiya K, Kawabata T, Tsuyuki T. 2005. A pre-operational variational data assimilation system for a non-hydrostatic model at the Japan meteorological agency: Formulation and preliminary results. *Q. J. R. Meteorol. Soc.* **131**: 3465–3475. DOI: 10.1256/qj.05.132

Janiskova M, Morcrette JJ. 2005. Investigation of the sensitivity of the ECMWF radiation scheme to input parameters using the adjoint technique. *Q. J. R. Meteorol. Soc.* **131**: 1975–1995. DOI: 10.1256/qj.04.183

Janiskova M, Mahfouf JF, Morcrette JJ, Chevallier F. 2002. Linearized radiation and cloud schemes in the ECMWF model: Development and evaluation. *Q. J. R. Meteorol. Soc.* **128**: 1505–1527. DOI: 10.1002/qj.200212858306

Kalnay E. 2003. *Atmospheric modelling, data assimilation and predictability*. Cambridge Univ. Press: Cambridge; ISBN 9780511076275. http://cdsweb.cern.ch/record/992314

Kalnay E, Li H, Miyoshi T, Yang SC, Ballbrera-Poy J. 2007. 4-D-Var or ensemble Kalman filter? *Tellus A* **59**: 758–773.

Karmitsa N. 2007. '*LMBM – FORTRAN* subroutines for Large-Scale nonsmooth minimization: User's manual', TUCS Technical Report 856. Turku Centre for Computer Science: Turku, Finland.

Karmitsa N, Bagirov A, Makela M. 2009. '*Empirical and theoretical comparisons of several nonsmooth minimization methods and software*,' Technical Report 959. Turku Centre for Computer Science: Turku, Finland. http://napsu.karmitsa.fi/publications/comparing.pdf

Lemarechal C. 1975. An extension of Davidon methods to non-differentiable problems. *Nondifferentiable Optimization, Mathematical Programming Studies* **3**: 95–109, DOI: 10.1007/BFb0120700.

Levy G, Coon M, Nguyen G, Sulsky D. 2010. Physically-based data assimilation. *Geoscientific Model Development Discussions* **3**: 517–540.

Lewis AS, Overton ML. 2008a. '*Behavior of BFGS with an exact line search on nonsmooth examples*,' Technical Report. Optimization Online. Mathematical Optimization Society: Philadelphia, PA http://www.optimization-online.org/DB_FILE/2008/12/2173.pdf

Lewis AS, Overton ML. 2008b. '*Nonsmooth optimization via BFGS*', Technical Report. Optimization Online. Mathematical Optimization Society: Philadelphia, PA http://www.optimization-online.org/DB_FILE/2008/12/2172.pdf

Liu DC, Nocedal J. 1989. On the limited memory BFGS method for large scale optimization. *Math. Program.* **45**: 503–528. http://portal.acm.org/citation.cfm?id=83726

Lorenc AC. 1986. Analysis methods for numerical weather prediction. *Q. J. R. Meteorol. Soc.* **112**: 1177–1194. DOI: 10.1002/qj.49711247414

Makela MM, Neittaanmaki P. 1992. *Nonsmooth optimization: analysis and algorithms with applications to optimal control*. World Scientific: Singapore; ISBN 9810207735.

Nocedal J, Wright S. 2006. *Numerical optimization*, 2nd edn. Springer: Berlin; ISBN 0387303030.

Skajaa A. 2010. '*Limited memory BFGS for nonsmooth optimization*', MS thesis, Courant Institute of Mathematical Science. New York University: New York, NY. http://www.cs.nyu.edu/overton/mstheses/skajaa/msthesis.pdf

Tippett MK, Anderson JL, Bishop CH, Hamill TM, Whitaker JS. 2003. Ensemble square root filters. *Mon. Weather Rev.* **131**: 1485–1490. DOI: 10.1175/1520-0493(2003)131⟨1485:ESRF⟩2.0.CO;2

Uzunoglu B, Fletcher SJ, Zupanski M, Navon IM. 2007. Adaptive ensemble reduction and inflation. *Q. J. R. Meteorol. Soc.* **133**: 1281–1294.

van Leeuwen PJ. 2001. An ensemble smoother with error estimates. *Mon. Weather Rev.* **129**: 709–728.

Wang Z, Navon IM, Dimet FX, Zou XL. 1992. The second order adjoint analysis: Theory and applications. *Meteorol. Atmos. Phys.* **50**: 3–20. DOI: 10.1007/BF01025501

Zhang S, Zou XL, Ahlquist JE, Navon IM, Sela JG. 2000. Use of differentiable and nondifferentiable optimization algorithms for variational data assimilation with discontinuous cost functions. *Mon. Weather Rev.* **128**: 4031–4044.

Zhang S, Zou XL, Ahlquist JE. 2001. Examination of numerical results from tangent linear and adjoint of discontinuous nonlinear models. *Mon. Weather Rev.* **129**: 2791–2804

Zhu J, Kamachi M, Zhou G. 2002. Nonsmooth optimization approaches to VDA of models with on/off parametrizations: Theoretical issues. *Adv. Atmos. Sci.* **19**: 405–424. DOI: 10.1007/s00376-002-0075-z

Zou XL, Navon IM. 1996. The linearization and adjoint of radiation transfer processes in the NMC spectral model part I: solar radiative transfer. *Meteorol. Atmos. Phys.* **58**: 193–203. DOI: 10.1007/BF01027565

Zou X, Navon IM, Le Dimet FX. 1992. Incomplete observations and control of gravity waves in variational data assimilation. *Tellus A* **44**: 273–296. DOI: 10.1034/j.1600-0870.1992.t01-2-00001.x

Zupanski D. 1993. The effects of discontinuities in the Betts–Miller cumulus convection scheme on four-dimensional variational data assimilation. *Tellus A* **45**: 511–524. DOI: 10.1034/j.1600-0870.1993.00013.x

Zupanski M. 2005. Maximum likelihood ensemble filter: theoretical aspects. *Mon. Weather Rev.* **133**: 1710–1726.

Zupanski D, Mesinger F. 1995. Four-dimensional variational assimilation of precipitation data. *Mon. Weather Rev.* **123**: 1112–1127.

Zupanski D, Zupanski M. 2006. Model error estimation employing an ensemble data assimilation approach. *Mon. Weather Rev.* **134**: 1337–1354

Zupanski M, Fletcher SJ, Navon IM, Uzunoglu B, Heikes RP, Randall DA, Ringler TD, Daescu D. 2006. Initiation of ensemble data assimilation. *Tellus A* **58**: 159–170. DOI: 10.1111/j.1600-0870.2006.00173.x

Zupanski M, Navon IM, Zupanski D. 2008. The maximum likelihood ensemble filter as a non-differentiable minimization algorithm. *Q. J. R. Meteorol. Soc.* **134**: 1039–1050.

