

Advanced Data Assimilation for Strongly Nonlinear Dynamics

GEIR EVENSEN

Nansen Environmental and Remote Sensing Center, Bergen, Norway

(Manuscript received 26 February 1996, in final form 18 June 1996)

ABSTRACT

Advanced data assimilation methods become extremely complicated and challenging when used with strongly nonlinear models. Several previous works have reported various problems when applying existing popular data assimilation techniques with strongly nonlinear dynamics. Common for these techniques is that they can all be considered as extensions to methods that have proved to work well with linear dynamics.

This paper examines the properties of three advanced data assimilation methods when used with the highly nonlinear Lorenz equations. The ensemble Kalman filter is used for sequential data assimilation and the recently proposed ensemble smoother method and a gradient descent method are used to minimize two different weak constraint formulations.

The problems associated with the use of an approximate tangent linear model when solving the Euler–Lagrange equations, or when the extended Kalman filter is used, are eliminated when using these methods. All three methods give reasonable consistent results with the data coverage and quality of measurements that are used here and overcome the traditional problems reported in many of the previous papers involving data assimilation with highly nonlinear dynamics.

1. Introduction

The celebrated Lorenz model has been the subject of extensive studies motivated by its chaotic and strongly nonlinear nature. In the field of data assimilation, the model has served as a testbed for examining the properties of various data assimilation methods when used with strongly nonlinear dynamics. The results have been used to suggest properties and possibilities of the methods for applications with oceanic and atmospheric models that may also be strongly nonlinear and chaotic.

In Gauthier (1992) the so-called adjoint method, which solves a variational minimization problem where the model acts as a strong constraint, was tested with the Lorenz model. It was found that in a case where the model did not undergo a transition, the cost function was relatively well behaved with respect to perturbations in the control variables—that is, the initial conditions for the three model variables. However, in a case where the model did undergo transitions a very strong sensitivity of the cost function with respect to perturbations in the initial conditions was observed—that is, the value of the cost function depended strongly and nonlinearly on the control variables and contained local minima.

Miller et al. (1994) reexamined the adjoint method with the Lorenz model and proved that this behavior of the cost function was strongly dependent on the length of the assimilation time interval. Gauthier (1992) presented an example with observations in the time interval $t \in [0, T]$ with $T = 8$, and two cases were compared, one with and one without transitions in the reference case, while Miller et al. (1994) examined three cases with chaotic behavior and $T = 8$, $T = 10$, and $T = 15$. These examples proved that the cost function became increasingly more sensitive with respect to small perturbations in the initial conditions as T was increased.

Miller et al. (1994) also gave a comprehensive discussion on applications of the extended Kalman filter with the Lorenz model. They found that the statistical linearization used in the extended Kalman filter, when deriving the error covariance evolution equation, resulted in a too simplified closure. The estimated solution was unreliable beyond $t = 11$. This was essentially explained by a poor prediction of error covariances resulting in insufficient gain because of a decaying mode that reflects the stability of the attractor. A generalization of the extended Kalman filter, where third- and fourth-order moments and evolution equations for these were included, was also examined. It was shown that this more sophisticated closure scheme provided a more consistent evolution of error statistics, which also resulted in sufficient gain to keep the estimate on track. Unfortunately, such an approach is not practical for a high-dimensional ocean or atmospheric model, since the fourth-order moment requires storage of N^4 elements (N

Corresponding author address: Geir Evensen, Nansen Environmental and Remote Sensing Center, Edvard Griegsvei 3a, N-5037 Solheimsviken, Norway.
E-mail: Geir.Evensen@nrsc.no

is the number of state variables). This can be compared with the second-order moment used in the extended Kalman filter where N^2 elements must be stored.

Other problems related to error variance prediction have also been observed with more realistic applications of the extended Kalman filter, for example, Evensen (1992) and Gauthier et al. (1993), who used the extended Kalman filter with quasigeostrophic models. Evensen (1992) pointed out that by evaluating the model operator at an unstable, say sheared, background flow this resulted in unbounded error variance growth. In Bouttier (1994) it was also concluded that unbounded error variance growth predicted by the error covariance equation in the Southern Hemisphere has to be limited by a representation of error saturation to account for nonlinearities in the model. So, in general, a more consistent closure is needed in the error covariance equation.

In this paper three recently developed data assimilation methods will be examined and intercompared in an example with the Lorenz equations. The first method is the gradient descent method, proposed by Evensen and Fario (1997), which minimizes a weak constraint variational formulation. The gradient descent method requires no forward or backward model integrations since the gradient is calculated with respect to the full state in space and time, and a new candidate for the solution is substituted in every iteration. This method will be briefly reviewed in section 2b.

Another weak constraint smoother method was recently proposed by van Leeuwen and Evensen (1996). The method applies ensemble integrations to represent the density for the model evolution in space and time and a variance minimizing estimate can then be calculated. The ensemble smoother method will be further discussed in section 2c.

The ensemble Kalman filter, which is a sequential method, was proposed by Evensen (1994a, 1994b) and used in a realistic application by Evensen and van Leeuwen (1996). The ensemble Kalman filter is based on a Monte Carlo approach for solving Kolmogorov's equation, which is the equation for evolution of the probability density function for the error statistics. For this method there is no need for any closure approximation. The ensemble Kalman filter will be further discussed in section 2d.

In the next section, the general formulation of the data assimilation problem and the methods used to solve it are discussed. An application of the three data assimilation methods is discussed in section 3 and a general discussion is given in section 4.

2. Formulation of the inverse problem

Now, the formulation of two weak constraint inverse problems for the Lorenz equations are presented. Both formulations define the maximum-likelihood estimator for the weak constraint problem as long as Gaussian prior error statistics can be assumed. Two vastly dif-

ferent solution methods are presented for the two formulations and will provide, respectively, the maximum-likelihood estimate and a variance minimizing estimate as the solution of the inverse problem. In addition, a sequential method, the ensemble Kalman filter, is briefly discussed.

a. Model equations

The Lorenz model consists of a system of three coupled and nonlinear ordinary differential equations (Lorenz 1963),

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x) + q^x, \\ \frac{dy}{dt} &= \rho x - y - xz + q^y, \\ \frac{dz}{dt} &= xy - \beta z + q^z.\end{aligned}\quad (1)$$

Here, $x(t)$, $y(t)$, and $z(t)$ are the dependent variables, and we have chosen the following commonly used values for the parameters in the equation; $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$. The terms $q^x(t)$, $q^y(t)$, and $q^z(t)$ are assumed to represent the unknown model errors. Initial conditions for the model are given as

$$\begin{aligned}x(0) &= x_0 + a^x, \\ y(0) &= y_0 + a^y, \\ z(0) &= z_0 + a^z,\end{aligned}\quad (2)$$

where x_0 , y_0 , and z_0 are the first-guess values of the initial conditions and the terms a^x , a^y , and a^z represent the errors in the first-guess initial conditions. If all the error terms were known or equal to zero, these equations would formulate a well-posed problem having a unique solution in a mathematical sense.

Now a set of measurements, $\mathbf{d} \in \mathfrak{R}^M$, of the true solution are assumed given and linearly related to the model variables by the measurement equation

$$\mathbf{d} = \mathcal{L}[x, y, z] + \boldsymbol{\epsilon}, \quad (3)$$

where $\mathcal{L} \in \mathfrak{R}^M$ is a linear measurement functional, $\boldsymbol{\epsilon} \in \mathfrak{R}^M$ is a vector of measurement errors, and M is the number of measurements.

b. Variational formulation

When the measurement equation is added to the system of equations, this results in an overdetermined problem and no solution can be found in general. However, by allowing the dynamics, the initial conditions, and the measurements to contain errors, a solution can be found that minimizes these error terms in a weighted least squares sense, for example, by minimizing the following variational integral,

$$\begin{aligned}
 J[x, y, z] = & \int_0^T dt_1 \int_0^T dt_2 \mathbf{q}(t_1)^T \mathbf{W}_{qq}(t_1, t_2) \mathbf{q}(t_2) \\
 & + \mathbf{a}^T \mathbf{W}_{aa} \mathbf{a} + \boldsymbol{\epsilon}^T \mathbf{W} \boldsymbol{\epsilon},
 \end{aligned} \tag{4}$$

where we define the error vectors $\mathbf{q}(t)^T = [q^x(t), q^y(t), q^z(t)]$ and $\mathbf{a}^T = (a^x, a^y, a^z)$. Here the weights $\mathbf{W}_{qq}(t_1, t_2)$ and \mathbf{W}_{aa} are 3×3 weight matrices, which are optimally defined as the inverses of the model error covariance matrix and the error covariance matrix for the initial conditions, respectively. The $M \times M$ matrix \mathbf{w} is the inverse of the measurement error covariance matrix.

Note that other estimators than least squares could be defined. However, the least squares formulation is attractive for several reasons. If the unknown errors are Gaussian, that is, completely explained by the two first statistical moments, mean and covariance, then minimizing (4) is equivalent to maximizing the probability density function,

$$P_j(x, y, z) = \frac{1}{Z_T} \exp(-J[x, y, z]). \tag{5}$$

Thus, the minimum of (4) is also the maximum-likelihood estimate as long as Gaussian prior error statistics are assumed. Further, when working with methods that involve the Euler–Lagrange equations these are readily derived and the derivatives of the penalty function exist everywhere.

A very simple approach for minimizing the penalty function (4) is to use a gradient descent algorithm as was proposed by Evensen and Fario (1997). Let first the model variables $x(t)$, $y(t)$, and $z(t)$ be represented on a numerical grid in time and stored in the state vectors \mathbf{x} , \mathbf{y} , and \mathbf{z} , all belonging to \mathfrak{R}^N , where N is the number of grid points in time.

The gradient of a discrete representation of $J[x, y, z]$, that is, $J[\mathbf{x}, \mathbf{y}, \mathbf{z}]$, with respect to the full state vector $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ in time is easily derived. When the gradient $\nabla_{\mathbf{x}, \mathbf{y}, \mathbf{z}} J$ is known it can be used in a descent algorithm to search for the minimizing solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}})$.

Since there is no integration of the model equations required in the gradient descent method used in this paper, very simple numerical discretizations based on second-order centered differences for the time derivatives will be used; that is,

$$\begin{aligned}
 \frac{x_{n+1} - x_{n-1}}{2\Delta t} &= \sigma(y_n - x_n) + q_n^x, \\
 \frac{y_{n+1} - y_{n-1}}{2\Delta t} &= \rho x_n - y_n - x_n z_n + q_n^y, \\
 \frac{z_{n+1} - z_{n-1}}{2\Delta t} &= x_n y_n - \beta z_n + q_n^z,
 \end{aligned} \tag{6}$$

where $n = 2: N - 1$ is the time-step index, with N the total number of time steps. In the ensemble integrations discussed in the next sections, and for the generation of the reference case, a standard high-order ordinary dif-

ferential equation solver has been used. The discrete model error is now written as $\mathbf{q}_n^T = (q_n^x, q_n^y, q_n^z)$. Further, as was also done by Evensen and Fario (1997), it is assumed for convenience that the model weight can be written as

$$\mathbf{W}_{qq}(t_1, t_2) = \mathbf{W}_{qq} \delta(t_1 - t_2), \tag{7}$$

where \mathbf{W}_{qq} is a constant 3×3 matrix and the time regularization, removed in this expression, will now be accounted for by a smoothing term acting on the inverse estimate.

The penalty function then becomes

$$\begin{aligned}
 J[\mathbf{x}, \mathbf{y}, \mathbf{z}] = & \Delta t \sum_{n=1}^N \mathbf{q}_n^T \hat{\mathbf{W}}_{qq} \mathbf{q}_n + \mathbf{a}^T \mathbf{W}_{aa} \mathbf{a} + \boldsymbol{\epsilon}^T \mathbf{W} \boldsymbol{\epsilon} \\
 & + \Delta t \sum_{n=1}^N \eta_n^T \mathbf{W}_{\eta\eta} \eta_n,
 \end{aligned} \tag{8}$$

where \mathbf{q}_n is defined in (6) and η_n is just the discrete centered second derivative acting on the inverse estimate at time step n . Note that \mathbf{q}_1 , \mathbf{q}_N , η_1 , and η_N use second-order one-sided difference formulas. It would have been more consistent to actually smooth the model errors instead of the inverse estimate, since it can be shown that such a smoothing constraint used together with the penalty term for the model errors would define a norm, and there is a unique correspondence between such a smoothing norm and a covariance matrix, as shown by McIntosh (1990). On the other hand will the smoothing term as included here improve the conditioning of the method since only smooth functions are searched for. The initial conditions and the measurements are included as before except that the measurement operator now must be considered as a matrix multiplied with the state vector consisting of all the discrete elements of \mathbf{x} , \mathbf{y} , and \mathbf{z} .

The required storage for the gradient descent method is of order the size of the state vector in space and time. Given a first-guess estimate, the gradient of the cost function is readily evaluated, and a new state estimate can be found.

When using nonlinear dynamics, the penalty function is clearly not convex in general due to the first term in (4) containing the model residuals. However, both the measurement penalty term and the smoothing norm will give a quadratic contribution to the penalty function, and if the weights \mathbf{w} and $\mathbf{W}_{\eta\eta}$ are large enough compared to the dynamical weight $\hat{\mathbf{W}}_{qq}$, one can expect a nearly quadratic penalty function. On the contrary, if the model residuals are the dominating terms in the penalty function, clearly a pure descent algorithm may get trapped in eventual local minima and the solution found may depend on the first guess in the iteration. For further details concerning the implementation of this algorithm, see Evensen and Fario (1997).

c. Probabilistic formulation

A different but still mathematically consistent formulation for the inverse problem can be given in terms of probability density functions. Such an approach was discussed in detail by van Leeuwen and Evensen (1996). Let first the model variables $x(t)$, $y(t)$, and $z(t)$ be represented on a numerical grid in time and stored in the state vector $\boldsymbol{\psi} \in \mathbb{R}^{3n}$, where n is the number of grid points in time. The determination of the generalized inverse is then considered as the estimation of the unknown true model variables $\boldsymbol{\psi}$ given the data and the model estimates with information about their prior error statistics. A Bayesian estimation problem can be formulated where the unknown $\boldsymbol{\psi}$ is viewed as the value of a random variable $\boldsymbol{\psi}'$. The probability density of the data \mathbf{d} is interpreted as the conditional distribution density $f(\mathbf{d} | \boldsymbol{\psi})$ of \mathbf{d}' assuming $\boldsymbol{\psi}' = \boldsymbol{\psi}$. The pure model is regarded as a priori information, and it is used to assign a density $f(\boldsymbol{\psi})$ to the random variable $\boldsymbol{\psi}'$. Using the definition of a conditional probability density we can derive the probability density of $\boldsymbol{\psi}$ given the data

$$f(\boldsymbol{\psi} | \mathbf{d}) = \frac{f(\mathbf{d} | \boldsymbol{\psi})f(\boldsymbol{\psi})}{f(\mathbf{d})}, \tag{9}$$

where the denominator can be rewritten as

$$f(\mathbf{d}) = \int f(\mathbf{d} | \boldsymbol{\psi})f(\boldsymbol{\psi}) d\boldsymbol{\psi}. \tag{10}$$

The joint probability density of the model evolution and the data is then given by

$$f(\boldsymbol{\psi} | \mathbf{d}) = \frac{f(\mathbf{d} | \boldsymbol{\psi})f(\boldsymbol{\psi})}{\int f(\mathbf{d} | \boldsymbol{\psi})f(\boldsymbol{\psi}) d\boldsymbol{\psi}}. \tag{11}$$

Thus, the probability density of the data given a model evolution, $f(\mathbf{d} | \boldsymbol{\psi})$, and the probability density of the model evolution, $f(\boldsymbol{\psi})$, must be known. The former is usually assumed to be known, for instance, a Gaussian. The value $\hat{\boldsymbol{\psi}}$ of $\boldsymbol{\psi}$ that maximizes $f(\boldsymbol{\psi} | \mathbf{d})$ is the maximum-likelihood estimate of $\boldsymbol{\psi}$.

For the Lorenz model, one has to specify initial conditions with their respective probability density. The probability $f(\boldsymbol{\psi})$ used above should therefore be written as $f(\boldsymbol{\psi} | \boldsymbol{\psi}_0)f(\boldsymbol{\psi}_0)$, and accordingly the probability density for the measurements should be $f(\mathbf{d} | \boldsymbol{\psi}, \boldsymbol{\psi}_0)$. Thus, the maximum-likelihood estimator maximizes the probability density

$$f(\boldsymbol{\psi}, \boldsymbol{\psi}_0 | \mathbf{d}) = Af(\mathbf{d} | \boldsymbol{\psi}, \boldsymbol{\psi}_0)f(\boldsymbol{\psi} | \boldsymbol{\psi}_0)f(\boldsymbol{\psi}_0), \tag{12}$$

or rather the log-likelihood function

$$\begin{aligned} \log f(\boldsymbol{\psi}, \boldsymbol{\psi}_0 | \mathbf{d}) &= \log f(\mathbf{d} | \boldsymbol{\psi}, \boldsymbol{\psi}_0) + \log f(\boldsymbol{\psi} | \boldsymbol{\psi}_0) \\ &\quad + \log f(\boldsymbol{\psi}_0) + \log A, \end{aligned} \tag{13}$$

where A is a constant arising from the denominator in Eq. (11).

In van Leeuwen and Evensen (1996) it was shown that if all prior error distributions are Gaussian, the max-

imization of the joint probability density (13) becomes equal to the minimization of the variational integral (4). A quadratic penalty function like (4) implicitly assumes that the errors are Gaussian. If this assumption fails to be true, the penalty function will no longer define the maximum-likelihood estimator, but it can still be used as a variance-minimizing estimator.

If one assumes that the model equations describe a first-order autoregressive, or Markov process—that is, the model is forced randomly as

$$d\boldsymbol{\psi} = \mathbf{g}(\boldsymbol{\psi})dt + d\boldsymbol{\beta}, \tag{14}$$

where $d\boldsymbol{\beta}$ are random increments with known covariance \mathbf{Q} and zero mean, the probability density, $f(\boldsymbol{\psi})$, for the model solution can be determined by solving the Kolmogorov equation,

$$\frac{\partial f}{\partial t} + \sum_{i=1}^n \frac{\partial g_i f}{\partial \psi_i} = \sum_{i,j=1}^n \frac{Q_{ij}}{2} \frac{\partial^2 f}{\partial \psi_i \partial \psi_j}. \tag{15}$$

A derivation of this equation, which is the fundamental equation for evolution of error statistics, can be found in Jazwinski (1970). The probability density function represents the density of an infinite ensemble of possible model states, each having an associated probability number. The width of the probability density function corresponds to the variance of the ensemble and represents the errors in the predicted solution.

The probability density for the model state has a huge number of variables, so it is computationally not feasible for real oceanographic or meteorologic applications to determine its evolution. An alternative is to solve (15) using an ensemble integration as discussed by Evensen (1994b), but to construct the density from the ensemble members is again not feasible for high dimensional state vectors. However, the first few moments of $f(\boldsymbol{\psi})$ are easily calculated from the ensemble.

Given the prior distribution $f(\boldsymbol{\psi} | \boldsymbol{\psi}_0)f(\boldsymbol{\psi}_0)$ for the model and the distribution for the data one can define a linear unbiased variance minimizing estimator by solving

$$\hat{\boldsymbol{\psi}} = \boldsymbol{\psi}_F + \mathbf{r}^T \mathbf{b}, \tag{16}$$

with

$$(\mathbf{R} + \mathbf{w}^{-1})\mathbf{b} = \mathbf{d} - \mathbf{L}\boldsymbol{\psi}_F, \tag{17}$$

where the matrix of representers or influence functions, $\mathbf{r} = \mathbf{L}\mathbf{Q}_{\boldsymbol{\psi}_F}$, and the representer matrix, $\mathbf{R} = \mathbf{L}\mathbf{Q}_{\boldsymbol{\psi}_F}\mathbf{L}^T$, can both be calculated from an ensemble of model solutions. Here, $\mathbf{Q}_{\boldsymbol{\psi}_F} \in \mathbb{R}^{3N \times 3N}$ is defined as the covariance matrix for the state variables including the time dimension. The discrete representation of the measurement functional \mathcal{L} —that is, the measurement matrix—is denoted \mathbf{L} : The first-guess estimate, $\boldsymbol{\psi}_F$, will be the mean of the ensemble, thus $\mathbf{Q}_{\boldsymbol{\psi}_F}$ can be interpreted as the error covariance of the first guess. Because of the nonlinearities in the model, this estimator is not the maximum-likelihood estimator, even if all the prior error statistics are Gaussian. However, it can be interpreted as the best linear

unbiased variance minimizing estimator or smoother for nonlinear model dynamics.

This method, the ensemble smoother, is similar to the analysis method used in the ensemble Kalman filter (Evensen 1994b), where a variance-minimizing estimate is calculated at measurement times based on the current ensemble statistics except that the smoother estimate is calculated over the whole space and time domain. Note that this ensemble smoother is as simple to calculate as the ensemble Kalman filter and no fields as function of both space and time need to be stored. The computational load is similar to that of the ensemble Kalman filter except that the ensemble members must be stored at diagnostic output times. The details for construction of the smoother solution using ensemble statistics were discussed in van Leeuwen and Evensen (1996).

Finally, it should be pointed out that it was shown by van Leeuwen and Evensen (1996) that if the model dynamics are linear then the maximum-likelihood estimate and the variance minimizing estimate are identical.

d. A sequential method: The ensemble Kalman filter

The ensemble Kalman filter is a sequential data assimilation method where the error statistics are predicted by solving the Kolmogorov's Eq. (15) using a Monte Carlo or ensemble integration. The method was originally proposed by Evensen (1994b) and has been further applied and discussed in Evensen (1994a), Evensen and van Leeuwen (1996), and van Leeuwen and Evensen (1996).

By integrating an ensemble of model states forward in time it is possible to calculate statistical moments like mean and error covariances whenever such information is required. Thus, all the statistical information about the predicted model state that is required at analysis times is contained in the ensemble.

In Evensen (1994b) an analysis scheme was proposed where the traditional update equation used in the Kalman filter is applied, except that the gain is calculated from the error covariances provided by the ensemble. Thus, Eqs. (16) and (17) are used for the analysis with the modification that the error covariance matrix \mathbf{Q}_{obs} is now local in time. It was also illustrated that a new ensemble with error statistics representing the analyzed state could be generated by updating each ensemble member individually using the same analysis equation. This analysis procedure can be characterized as the optimal variance-minimizing method. An inherent assumption is that the error statistics are Gaussian with vanishing higher-order statistical moments. This is, in general, not true for nonlinear dynamics where the probability density function may be far from Gaussian and higher-order moments may contribute significantly [see Miller (1994) for example].

The ensemble Kalman filter avoids many of the problems associated with the traditional extended Kalman

filter, for example, there is no closure problem as is introduced in the extended Kalman filter by neglecting contributions from higher-order statistical moments in the error covariance evolution equation. It can also be computed at a much lower numerical cost, since only a few hundred model states may be sufficient for reasonable statistical convergence. When the size N_{ens} of the ensemble increases, the errors in the solution for the probability density will approach zero at a rate proportional to $N_{\text{ens}}^{-1/2}$. For practical ensemble sizes, say $O(100)$, the errors will be dominated by statistical noise, not by closure problems or unbounded error variance growth as have been observed in the extended Kalman filter (see Evensen 1992, 1994b). A similar statistical convergence can also be expected in the ensemble smoother.

3. Examples

For all the cases to be discussed, the initial condition for the reference case is given by $(x_0, y_0, z_0) = (1.508870, -1.531271, 25.46091)$ and the time interval is $t \in [0, 40]$. The observations and initial conditions are simulated by adding normal distributed noise with zero mean and variance equal to 2.0 to the reference solution. The initial conditions used are also assumed to have the same variance as the observations. These are the same values that were used in Miller et al. (1994) and Evensen and Fario (1997).

The model error covariance matrix $\hat{\mathbf{Q}}_{qq}$ from Evensen and Fario (1997),

$$\hat{\mathbf{Q}}_{qq} = \begin{bmatrix} 0.1491 & 0.1505 & 0.0007 \\ 0.1505 & 0.9048 & 0.0014 \\ 0.0007 & 0.0014 & 0.9180 \end{bmatrix}, \quad (18)$$

corresponding to a time step equal to $\Delta t = 0.0167$ is used for the gradient descent method. In the ensemble-based methods a stochastic forcing term with similar error statistics is applied. For the time smoothing constraint used in the gradient descent method, the same value as was found by Evensen and Fario (1997) is used; that is, the smoothing weight matrix is chosen to be diagonal and given by $\mathbf{W}_{\eta\eta} = 10^{-3}\mathbf{I}$.

a. Experiment A

The three methods discussed above will now be examined and compared in an experiment where the distance between the measurements is $\Delta t_{\text{obs}} = 0.25$, which is the same as was used in Miller et al. (1994). Thus, it is possible to compare the results presented here with those presented in Miller et al. (1994) using the extended Kalman filter and a strong constraint variational method.

1) GRADIENT DESCENT SOLUTION

The maximum-likelihood estimate for a weak constraint inverse problem where Gaussian error statistics




FIG. 1. Experiment A (gradient descent): The inverse estimate for x (upper) and the corresponding error variance estimate (lower) vs time. The estimated solution is given by the solid line. The dashed line (hardly distinguishable from the solid line) is the true reference solution, and the diamonds show the simulated observations. The same line types will be used also in the following figures.

have been assumed can be solved using the gradient descent method discussed in section 2b.

The first-guess solution was initially chosen as the mean of the reference solution—that is, about $(0, 0, 23)$. It was found that using the gradient descent method, the global minimum was always found as long as the measurement density and quality was reasonably good. However, when the measurement errors became larger, or if a low number of measurements were used, the gradient descent method often converged to a local minimum. There also seemed to be a possibility for a local minima close to the zero solution where both the dynamical penalty term and the smoothing penalty vanish. It is therefore not wise to use an estimate close to the zero solution as the first guess in the descent algorithm.

To reduce the probability of getting trapped in eventual local minima, an objective analysis estimate, con-

sistent with the measurements, was used as a first guess in the descent algorithm. It was calculated using a smoothing spline minimization algorithm, which is equivalent to objective analysis (McIntosh 1990). This was easily done by replacing the dynamical misfit term with a penalty on a first-guess estimate in the penalty function (8).

The gradient descent method was in the current example capable of finding the global minimum when starting from the objective analysis estimate. The minimizing solution for the x component is given in Fig. 1 together with the error variance estimate. It is amazing how close the inverse estimate is to the reference solution even with such large errors in the measurements. The quality of this inverse estimate is clearly superior to previous inverse calculations using the extended Kalman filter or a strong constraint formulation.




FIG. 2. Experiment A (ensemble smoother): The inverse estimate for x (upper) and the corresponding error variance estimate (lower).

The error estimate has been generated using a statistical sampling method (Evensen and Fario 1997). This approach utilizes the fact that the minimum solution can be interpreted as the maximum-likelihood estimate of a probability density function (5). By using a hybrid Monte Carlo method to generate a Markov chain that samples this function, a statistical variance estimate can be generated. Note that this method may be used to generate error estimates independently of the minimization technique used to solve the weak constraint problem. The estimated error variances for the x component are given in Fig. 1. The largest errors appear around the peaks of the solution.

2) ENSEMBLE SMOOTHER SOLUTION

In the calculation of the ensemble smoother estimate, an ensemble of 1000 members was used. The same simulation was rerun with various ensemble sizes and the

differences between the results were minor, for example, using 500 members gave essentially the same result as the 1000 member case. The ensemble smoother solution for the x component and its estimated error variance is given in Fig. 2. The estimate is not as close to the reference solution as was the case for the gradient descent method, but that can be expected since this is a variance minimizing estimate, and the posterior probability distribution will not be a Gaussian because of the nonlinearities in the model. A positive result is that the smoother estimate is not missing any of the transitions and the performance is only poor when it comes to reproducing the peaks of the reference solution. The error variance estimate is given in the lower plot of Fig. 2. In the ensemble smoother the posterior error variances can be easily calculated by performing an analysis for each of the ensemble members and then evaluating the variance of the new ensemble. Clearly, the error estimates are not large enough at the peaks where the




FIG. 3. Experiment A (ensemble Kalman filter): The inverse estimate for x (upper) and the corresponding error variance estimate (lower).

smoother performs poorly. This is again a result of neglecting the non-Gaussian contribution from the probability distribution for the model evolution. Thus, the method assumes the distribution is Gaussian and believes it is doing well. Otherwise the error estimate looks reasonable with minima at the measurement locations and maxima in between the measurements. Note again that if a linear model is used, the posterior density will be Gaussian and the ensemble smoother will, in the limit of an infinite ensemble size, provide the same solution as the gradient descent method.

3) ENSEMBLE KALMAN FILTER SOLUTION

Finally, the data assimilation estimate from the ensemble Kalman filter has been included for comparison, and the results are given in Fig. 3. Also here the ensemble consists of 1000 members. The ensemble Kalman filter seems to do a reasonably good job in tracking

the phase transitions and also in reproducing the correct amplitudes in the peaks of the solution. There are a few locations where the filter estimate starts diverging from the reference solution, for example, for $t = 26$ and $t = 35$. Note, however, that the ensemble Kalman filter recovers quickly and begins tracing the reference solution again. The error estimate given in the lower plot of Fig. 3 shows strong error growth at these particular locations and thus indicates that the ensemble is passing through a region in the state space, which is associated with strong instability.

The error estimates show the same behavior as was found by Miller et al. (1994) with very strong error growth when the model solution is passing through the unstable regions of the state space, and otherwise rather weak error variance growth in the more stable regions. Note, for example, the low error variance when $t \in [28, 34]$ corresponding to the oscillation of the solution around one of the attractors. A problem for the filter is

that if the reference solution had made a transition at, say $t = 30$, then the rather low variance estimate might not be sufficient to produce a large enough gain to pull the filter estimate out of the stable oscillation. This issue is further discussed in Miller et al. (1994).

The maybe surprising result is that the ensemble Kalman filter seems to perform better than the ensemble smoother. At least this is surprising based on linear theory, where one has learned that the Kalman smoother solution at the end of the time interval is identical to the Kalman filter solution, and the additional information introduced by propagating the contribution of future measurements backward in time further reduces the error variance compared to the filter solution. For a nonlinear model this will still be true if the extended Kalman filter and smoother are used. However, for a nonlinear model neither the extended Kalman filter or the extended Kalman smoother will be optimal since they are based on linearized dynamics for the error covariance equation. Thus, even if the extended Kalman filter and smoother can be derived from a variational formulation like (4), as a method for solving the Euler–Lagrange equations, these methods will not give the same result as a direct substitution method solving (4) directly. Note again that if the model dynamics are linear, the ensemble Kalman filter will give the same solution as the Kalman filter, and the ensemble smoother will give the same result as the Kalman smoother.

Finally, it should be pointed out that in the ensemble filter a variance minimizing analysis is calculated at measurement times. Thus, even if the ensemble certainly is non-Gaussian due to the forward integration of nonlinear model equations, only the Gaussian part of the distribution is used. This is in contrast to the work by Carter et al. (1996, personal communication), where the maximum-likelihood analysis is calculated by actually constructing the density function for the model evolution and then calculating the conditional density in terms of analytical functions. They found that this made a significant improvement on the analysis, however, it is still not clear how this approach can be used in a practical way for high-dimensional state spaces.

b. Experiment B

To examine the sensitivity of the three methods discussed in the previous section, with respect to measurement density, an additional experiment is now performed where the distance between the measurements is $\Delta t_{\text{obs}} = 0.5$.

1) GRADIENT DESCENT SOLUTION

The solution found when using the gradient descent method is given in Fig. 4. The lower data density causes the inverse estimate to miss three transitions, thus, the method converges to a local minimum. This implies that the lower data density reduces the quadratic contribution

from the measurement functional, compared to experiment A, and the nonlinear model dynamics now have a larger relative impact on the shape of the penalty function.

Interestingly, the error estimate is similar to the one found for the gradient descent solution in experiment A, and it is not at all capable of capturing the effect of the transitions that were missed. This is caused by an improper sampling of the penalty function, where all candidates are taken in the neighborhood of the estimate within the well of the local minimum. A proper sampling is very expensive and would of course also have found the global minimum of the penalty function. Such a Monte Carlo approach for minimizing the penalty function was discussed in Evensen and Fario (1997). Note also that the missed transitions correspond to regions of maximum variances from experiment A, which may suggest that these are sensitive regions also for the higher data density in experiment A. The positive result from this experiment is that the formulation and method used give a very accurate solution over most of the domain even if a few transitions are missed.

2) ENSEMBLE SMOOTHER SOLUTION

The ensemble smoother performs rather poorly with the lower data density, as shown in Fig. 5. This can be expected in such a strongly nonlinear case where the measurements do not properly resolve the characteristic timescale of the system. Note that the smoother solution consists of a first-guess estimate (mean of the freely evolving ensemble) plus a linear combination of time-dependent influence functions or representers, which are calculated from the ensemble statistics. Thus, the method becomes equivalent to a variance minimizing objective analysis method where the time dimension is included. The estimated error variances are significantly larger than in experiment A, which informs us about the poor quality of the solution, but still seems to underestimate the actual errors in the inverse estimate. The ensemble smoother seems to do a better job with weakly nonlinear dynamics or with higher measurement density, which is natural since it is a variance minimizing method.

3) ENSEMBLE KALMAN FILTER SOLUTION

The ensemble Kalman filter does a reasonably good job in tracking the reference solution with the lower data density, as can be seen in Fig. 6. The estimate is not as close to the reference solution as was the case for the gradient descent method, after all, this is a sub-optimal filter estimate. On the other hand, only one transition completely missed, compared to three when using the gradient descent method.

The error variance estimate is also rather consistent and is at least capable of capturing the missed transition at $t = 18$, although the error variances are also large at




FIG. 4. Experiment B (gradient descent): The inverse estimate for x (upper) and the corresponding error variance estimate (lower) vs time. The estimated solution is given by the solid line. The dashed line (hardly distinguishable from the solid line) is the true reference solution, and the diamonds show the simulated observations. The same line types will be used also in the following figures.

another three locations, where the first is associated with the divergence of the ensemble from the initial condition and the last is in a region, at $t = 34$ where the estimate clearly attempts to diverge from the reference solution but is pulled back by measurements. The reason for the peak at $t = 10$ is more unclear though.

4. Discussion

Two weak constraint variational methods and a sequential method have been compared in two data assimilation experiments with the Lorenz equations. The first weak constraint method is a gradient descent method, which is used to minimize the penalty function, and it can be shown that if the prior error statistics are Gaussian, then the minimizing solution becomes the maximum-likelihood estimate. The second method is the en-

semble smoother, which provides a variance-minimizing estimate by rejecting the non-Gaussian part of the density of the pure model evolution. In a case with reasonable data coverage, these two methods are both capable of reproducing all the state transitions in the reference solution although the maximum-likelihood estimate is superior in reproducing the peaks in the reference solution. Also the ensemble Kalman filter does a good job in tracking the reference solution. The filter estimate is actually better than the ensemble smoother estimate in reproducing the peaks of the reference solution, and with low data density the ensemble smoother gives a rather poor result.

These examples illustrate how complicated the data assimilation problem becomes when strongly nonlinear dynamics are used. Clearly, there are several solution methods for the inverse problem which for nonlinear




FIG. 5. Experiment B (ensemble smoother): The inverse estimate for x (upper) and the corresponding error variance estimate (lower).

models result in different solutions, while it can be formally shown that all these methods give identically the same solution when used with linear dynamics. Both based on the examples discussed in this paper and from the common understanding of statistical estimation theory, the maximum-likelihood estimate is normally recognized as the optimal estimate. Thus, a substitution method for minimizing (4), for example, the gradient descent method or the even more sophisticated simulated annealing method used by Bennett and Chua (1994), is probably the only practical approach that can be applied when solving for the maximum-likelihood estimate using strongly nonlinear dynamics. Alternatives may be either to construct the probability density for the model evolution following the lines of Carter et al. (1996, personal communication), or to use the representer method to iterate the Euler–Lagrange equations (Bennett 1992). However, the former of these is not yet

practical for high-dimensional state spaces and the latter has been applied only for weakly nonlinear dynamics, for example, Bennett and Thorburn (1992) and Bennett et al. (1993).

Using the ensemble Kalman filter we obtained good solutions using down to about 10 ensemble members, but this is probably caused by the fact that all three state variables are observed at measurement times, and the influence of covariances between variables then becomes less important. In Evensen (1994b), the sensitivity of the ensemble Kalman filter solution with respect to the size of the ensemble was further discussed, and experiments were successful for a multilayer QG model using 100 ensemble members. However, it is not yet clear if this number must be increased for applications with primitive equations models. The ensemble smoother may need a larger ensemble size since covariances also in time are required.




FIG. 6. Experiment B (ensemble Kalman filter): The inverse estimate for x (upper) and the corresponding error variance estimate (lower).

On the other hand, when the dynamics is only weakly nonlinear, several solution methods may be used. In addition to the gradient descent method, the representer method will probably converge, and the distribution for the model evolution will probably be close to Gaussian so that the ensemble smoother should provide good results, too. For primitive equation models the ensemble Kalman filter or the ensemble smoother may be the most practical alternatives because of the lesser storage requirements compared to the gradient descent method, and one could use an existing model with only limited additional coding for the analysis step to have a running data assimilation system. The gradient descent method is perhaps more practical for lower dimensional state spaces, for example, when simpler dynamical constraints are added to a penalty function in a generalized inverse formulation for inverting a dataset. Such an ex-

ample has been discussed by Zaron (1995), using time-independent dynamical constraints.

These examples also illustrate that by allowing the dynamics to contain errors, that is, the model is imposed as a weak constraint, the extreme sensitivity of the penalty function with respect to initial conditions seen in strong constraint applications is completely removed. Further, there are no limitations on the length of the assimilation interval.

When using a gradient descent method to solve a weak constraint problem the full model state in space and time must be stored simultaneously. Thus, the numerical load is mainly associated with the storage problem. If an efficient gradient descent type solution method is used, it is expected that the CPU requirements are lower than for the ensemble smoother, which requires the forward integration of an ensemble of model states.

On the other hand, the ensemble smoother does not require the storage of a model state in both space and time. Actually, except for the analysis step, the method requires only a sequence of independent forward integrations of the model.

Acknowledgments. This work was supported by the Norwegian Research Council.

REFERENCES

- Bennett, A. F., 1992: *Inverse Methods in Physical Oceanography*. Cambridge University Press, 346 pp.
- , and M. Thorburn, 1992: The generalized inverse of a nonlinear quasigeostrophic ocean circulation model. *J. Phys. Oceanogr.*, **22**, 213–230.
- , and B. S. Chua, 1994: Open-ocean modeling as an inverse problem: The primitive equations. *Mon. Wea. Rev.*, **122**, 1326–1336.
- , L. M. Leslie, C. R. Hagelberg, and P. E. Powers, 1993: Tropical cyclone prediction using a barotropic model initialized by a generalized inverse method. *Mon. Wea. Rev.*, **121**, 1714–1729.
- Bouttier, F., 1994: A dynamical estimation of forecast error covariances in an assimilation system. *Mon. Wea. Rev.*, **122**, 2376–2390.
- Evensen, G., 1992: Using the extended Kalman filter with a multilayer quasi-geostrophic ocean model. *J. Geophys. Res.*, **97**(C11), 17 905–17 924.
- , 1994a: Inverse methods and data assimilation in nonlinear ocean models. *Physica D*, **77**, 108–129.
- , 1994b: Sequential data assimilation with a nonlinear quasigeostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**(C5), 10 143–10 162.
- , and N. Fario, 1997: A weak constraint variational inverse for the Lorenz equations using substitution methods. *J. Meteor. Soc. Japan*, in press.
- , and P. J. van Leeuwen, 1996: Assimilation of Geosat altimeter data for the Agulhas current using the ensemble Kalman filter with a quasigeostrophic model. *Mon. Wea. Rev.*, **124**, 85–96.
- Gauthier, P., 1992: Chaos and quadri-dimensional data assimilation: A study based on the Lorenz model. *Tellus, Ser. A*, **44**, 2–17.
- , P. Courtier, and P. Moll, 1993: Assimilation of simulated wind lidar data with a Kalman filter. *Mon. Wea. Rev.*, **121**, 1803–1820.
- Jazwinski, A. H., 1970: *Stochastic Processes and Filtering Theory*. Academic, 376 pp.
- Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141.
- McIntosh, P. C., 1990: Oceanic data interpolation: Objective analysis and splines. *J. Geophys. Res.*, **95**(C8), 13 529–13 541.
- Miller, R. N., 1994: Perspectives on advanced data assimilation in strongly nonlinear systems. *Data Assimilation: Tools for Modelling the Ocean in a Global Change Perspective*, P. P. Brasseur and J. C. J. Nihoul, Eds., NATO ASI, Vol. I 19. Springer-Verlag, 195–216.
- , M. Ghil, and F. Gauthiez, 1994: Advanced data assimilation in strongly nonlinear dynamical systems. *J. Atmos. Sci.*, **51**, 1037–1056.
- van Leeuwen, P. J., and G. Evensen, 1996: Data assimilation and inverse methods in terms of a probabilistic formulation. *Mon. Wea. Rev.*, **124**, 2898–2913.
- Zaron, E., 1995: The time-averaged circulation of the North Pacific Ocean: An analysis based on inverse methods. Ph.D. dissertation, Oregon State University, 211 pp. [Available from College of Oceanic and Atmospheric Sciences, 104 Ocean Administrative Building, Corvallis OR 97331-5503.]