© Copyright 2008 American Meteorological Society (AMS). Permission to use figures, tables, and brief excerpts from this work in scientific and educational works is hereby granted provided that the source is acknowledged. Any use of material in this work that is determined to be “fair use” under Section 107 of the U.S. Copyright Act September 2010 or that satisfies the conditions specified in Section 108 of the U.S. Copyright Act (17 USC §108, as revised by P.L. 94-553) does not require the AMS’s permission. Republication, systematic reproduction, posting in electronic form, such as on a web site or in a searchable database, or other uses of this material, except as exempted by the above statement, requires written permission or a license from the AMS. Additional details are provided in the AMS Copyright Policy, available on the AMS Web site located at (<http://www.ametsoc.org/>) or from the AMS at 617-227-2425 or copyright@ametsoc.org.

Analysis of the Impact of Model Nonlinearities in Inverse Problem Solving

T. VUKICEVIC

ATOC, and CIRES, University of Colorado, Boulder, Colorado

D. POSSELT

ATS, Colorado State University, Fort Collins, Colorado

(Manuscript received 8 June 2007, in final form 30 January 2008)

ABSTRACT

In this study, the relationship between nonlinear model properties and inverse problem solutions is analyzed using a numerical technique based on the inverse problem theory formulated by Mosegaard and Tarantola. According to this theory, the inverse problem and solution are defined via convolution and conjunction of probability density functions (PDFs) that represent stochastic information obtained from the model, observations, and prior knowledge in a joint multidimensional space. This theory provides an explicit analysis of the nonlinear model function, together with information about uncertainties in the model, observations, and prior knowledge through construction of the joint probability density, from which marginal solution functions can then be evaluated. The numerical analysis technique derived from the theory computes the component PDFs in discretized form via a combination of function mapping on a discrete grid in the model and observation phase space and Monte Carlo sampling from known parametric distributions. The efficacy of the numerical analysis technique is demonstrated through its application to two well-known simplified models of atmospheric physics: damped oscillations and Lorenz's three-component model of dry cellular convection. The major findings of this study include the following: (i) Use of a nonmonotonic forward model in the inverse problem gives rise to the potential for a multimodal posterior PDF, the realization of which depends on the information content of the observations and on observation and model uncertainties. (ii) The cumulative effect of observations over time, space, or both could render the final posterior PDF unimodal, even with the nonmonotonic forward model. (iii) A greater number of independent observations are needed to constrain the solution in the case of a nonmonotonic nonlinear model than for a monotonic nonlinear or linear forward model for a given number of degrees of freedom in control parameter space. (iv) A nonlinear monotonic forward model gives rise to a skewed unimodal posterior PDF, implying a well-posed maximum likelihood inverse problem. (v) The presence of model error greatly increases the possibility of capturing multiple modes in the posterior PDF with the nonmonotonic nonlinear model. (vi) In the case of a nonlinear forward model, use of a Gaussian approximation for the prior update has a similar effect to an increase in model error, which indicates there is the potential to produce a biased mean central estimate even when observations and model are unbiased.

1. Introduction

The use of techniques for solving inverse problems is necessary in the earth system sciences because these techniques provide the only means for obtaining an objective representation of mutually consistent, quasi-continuous, spatially and temporally variant fields of physical quantities from a variety of environmental observations. In atmospheric and oceanic science applica-

tions, the solution to an inverse problem is typically referred to as state estimation or data assimilation because the problem is formulated as assimilation of measurements into environmental models for the purpose of estimating the state of the system (Kalnay 2003; Lewis et al. 2006; Bennett 1992). An elegant review of the motivation for and the benefits derived from data assimilation in the earth system is presented in O'Neill et al. (2004). Their review recognizes the special role of satellite remote sensing observations in modern environmental state analysis, and it is worth noting that methods for solving inverse problems have been used to retrieve environmental parameters from

Corresponding author address: T. Vukicevic, ATOC, University of Colorado, UCB 311, Boulder, CO 80309.
E-mail: tomlava.vukicevic@colorado.edu

these measurements for quite some time (Rodgers 2000).

Another important use of methods for solving inverse problems is in research designed to improve numerical models of the environment through objective estimation of model physical parameters in a way that is consistent with the observations, the model, and their uncertainties (Annan et al. 2005; Aksoy et al. 2006; Braswell et al. 2005; Vukicevic et al. 2001). In contrast to estimation of quasi-continuous environmental fields, the use of inverse methods is not explicitly necessary in this domain, but results indicate that improvements could be made in modeling via the objective correction that results from inverse problem solutions.

Because of broad interest in inverse problem solving in the earth system sciences, a large number of studies and several new books have been published in last decade on the subject of the theoretical background and technical implementation of the methods (Rodgers 2000; Kalnay 2003; Lewis et al. 2006; Evensen 2006 and references therein). Although the fundamental theory used in solving inverse problems has a long history of development outside of and prior to its modern implementation in the earth system sciences, applications within these sciences emphasize two of the most challenging aspects of the problem. These are (i) characterization of a large multivariate and multidimensional phase space and (ii) use of complex nonlinear models with poorly known uncertainties.

The problem of a large phase space has been addressed through use of data assimilation techniques that incorporate efficient computation of sensitivities of model simulated measurements to control parameters or state, together with advances in computing power. Contemporary data assimilation techniques belong to one of two classes: variational least squares and ensemble filters or smoothers. Variational least squares techniques are spatial variational techniques that include three- and four-dimensional variational data assimilation (3DVAR, 4DVAR), where 3D refers to three spatial dimensions and 4D includes a temporal dimension in addition to the spatial ones. Ensemble filters and smoothers are all techniques that base their theoretical formulation on the Kalman filter (KF) approach (Kalman 1960) and make use of an ensemble of model simulations. Excellent reviews of the techniques commonly used in the earth system sciences today are presented in Lewis et al. (2006) and Evensen (2006). In modern techniques, the effect of model nonlinearity on the inverse problem solution is treated primarily by allowing the nonlinear model to respond to a change in the controlling parameters or state, where the change is computed under criteria that include the linear model

assumption. This approach does not necessarily include the full effect of model nonlinearity on the inverse problem solution.

It is also worth mentioning that although each modern data assimilation technique requires computation of a change in computed observations with respect to a change in control parameters or initial conditions, the data assimilation studies referenced above typically do not provide a priori analysis of the validity of the assumptions that underlie these computations. Instead, the overall approach is most often justified a posteriori by evaluating the improvement in simulated observations relative to the actual observations. Additional indirect verification is provided by showing that the model forecast subsequent to the data assimilation is also improved. Although these positive results are certainly desirable and justify the effort put into development and application of a particular technique, they do not provide explicit proof that nonlinear model effects have been treated appropriately. It follows that it is not straightforward either to anticipate the skill of the technique in applications that employ a different nonlinear model or to assess possible causes of loss of skill when the technique does not perform well in a new application.

The discrepancy between nonlinear model effects and the criteria that define the change of state or parameters results primarily from assumptions made about the nature of the stochastic posterior solution (i.e., the solution of the inverse problem). The common assumption is that the posterior probability density function (PDF) is approximately Gaussian. It is well known, however, that stochastic estimation theory implies that these properties are not satisfied for a general nonlinear model (Jazwinski 1970; Tarantola 2005). Several recent studies (Kim et al. 2003; van Leeuwen 2003) have suggested alternative approaches for nonlinear systems with a large phase space that are not based on linear theory (i.e., based on the assumption that the posterior PDF is Gaussian). These studies proposed the use of particle filter techniques, which do not impose a specific parametric form on the posterior PDF (Liu 2001). Chin et al. (2007) have proposed an ensemble-based smoother, termed the backward sequential smoother (BSS), as yet another approach that does not require the posterior PDF to be Gaussian. Although these recent studies have addressed the need to consider non-Gaussian statistics in nonlinear inverse problem solving, they have not offered an explicit analysis of the relationship between model nonlinearity and the properties of the posterior PDF.

In this study, we illuminate the relationship between nonlinear model properties and the inverse problem

solution by analyzing the role of the nonlinear model in shaping the posterior PDF. This is done via numerical analysis of the properties of the posterior PDF. The analysis method is based on the general stochastic inverse problem theory as set forth in Mosegaard and Tarantola (2002, hereafter MT02). Similar to standard Bayesian statistical estimation theory (Jazwinski 1970; Evensen 2006; Bennett 1992) the formulation in MT02 treats each quantity that contributes to a total quantitative knowledge of modeled and observed states of the system as a stochastic quantity with an associated PDF. The theory used by MT02 is described in great detail in the book by Tarantola (2005). A unique feature of the MT02 approach is that contributions from the model and the associated modeling errors, the observation errors, and the prior information are each represented by a separate PDF which, when combined together by convolution and conjunction, result in a joint posterior PDF. The properties of the joint posterior PDF explicitly determine the type of inverse problem, such as whether or not the solution can be expected to be unimodal. The explicit separation of information into components provided by the model, observation, and prior estimate allows unique identification of the contribution of each to the posterior PDF.

In this study, we develop a numerical method of analysis based on the MT02 approach for the purpose of diagnosing the properties of the nonlinear inverse problem. The analysis is performed for two relatively simple nonlinear models: the damped oscillator model and Lorenz's three-component chaos model (Lorenz 1963). The remainder of this paper is organized as follows: The general formulation of the inverse problem theory following MT02 is presented in section 2a. In section 2b, it is shown that under the typical assumption that PDFs are Gaussian for the model error, observation errors, and the prior estimate, this theory results in the well-known formulation of the data assimilation problem that serves as the basis for modern data assimilation techniques. The method of numerical nonlinear analysis is described in section 3, and results of the analysis applied to the examples of damped oscillations and Lorenz's three-variable chaos model (Lorenz 1963) are presented in section 4. A brief summary and general conclusions are presented in section 5.

2. Inverse problem formulation

a. General theory

By the approach in MT02, three sources of discrete quantitative information about a system under study can be identified. These are the physical measurements, prior knowledge, and a model of the system of interest

that links observations to the desired parameters. Each piece of information resides in a corresponding space containing all possible values. The parameter space is assumed to be generally multidimensional and is spanned by points $m \in [(m_i^1, \dots, m_i^M), i = 1, I]$, where M is the phase dimension of the parameter space (i.e., the number of physical quantities) and I is a number of possible discrete realizations for any given resolution (i.e., the unit volume in the discretized space). The parameter space is transformed into a discrete measurement space by the model

$$y = f(m), \quad (2.1)$$

where f is the (nonlinear) model operator, typically referred to as the forward model. For convenience the parameter space is denoted Π . The measurement space as simulated by the model, denoted S , is spanned by points $y \in [(y_l^1, \dots, y_l^N), l = 1, L]$, where N is the phase dimension of the observation space (the number of different physical measurements) and L is a number of possible discrete realizations. Similar to the parameter space, the discretization assumes the adoption of a resolution. The joint space $\Omega = \Pi \times S$ is characterized by a joint probability density function $p_1(m, y)$ and is the space that contains all possible information available about the system, given the choice of model and measurement space (e.g., the observing network). The joint probability density on Ω provides a complete description of the modeling uncertainties and natural variability in the parameters.

The other piece of information about the system is contained in the actual physical measurements, which are assumed to be independent of the model. Let this information be contained in a space denoted Y . There is a joint probability density on the joint space $\Theta = Y \times \Pi$, denoted $p_2(m, y)$. The union of measurement spaces S and Y defines the total measurement space, which is denoted D . A new set of information about the system is obtained when the two joint probability densities $p_1(m, y)$ and $p_2(m, y)$ are combined by conjunction of PDFs to result in the joint probability density in $D \times \Pi$. The conjunction is written

$$p(m, y) = \frac{1}{\gamma} \frac{p_1(m, y)p_2(m, y)}{v(m, y)}, \quad (2.2)$$

where $\gamma = \int_{DM} [p_1(m, y)p_2(m, y)]/[v(m, y)]$ is the normalization constant and v is the homogenous probability density only consisting of the information that equal probabilities are assigned to equal volumes in the discrete space $D \times \Pi$.

The a posteriori probability density $p(m, y)$ on the joint space $D \times \Pi$ results from the combined joint probability distributions of information about the si-

multaneously modeled and measured system. The knowledge of this probability density is the most complete available quantitative knowledge about the system; by this property, the expression (2.2) defines the general inverse problem. Note that this is the most general statement of the inverse problem; we have made no assumptions about the form of the probabilities or the shape of the information space itself. The solution is the marginal PDF in the parameter space:

$$p_m(m) = \int_D p(m, y) dy. \quad (2.3)$$

To arrive at the solution (2.3), each of the PDFs on the right-hand side of (2.2) must be specified.

The specification of the joint probability density for the system given only the model, without the observations and prior knowledge, is as follows: First, by definition of conditional and marginal PDFs, $p_1(m, y)$ can be written

$$p_1(m, y) = p_1(y/m)v(m). \quad (2.4)$$

Second, the marginal probability density in the parameter space, in the absence of prior knowledge and the use of actual measurements, is assumed to be equivalent to the homogenous probability density of the parameters (i.e., the knowledge of their existence only). This PDF is simply constant for linear parameter spaces, which is a common assumption in geophysical models, but it should be noted that nonnegative physical parameters that are typically near zero require special attention. The conditional probability density $p_1(y/m)$ is composed of results of the forward model when it is applied over the space of the control parameters with the homogenous PDF without any knowledge of the measurements or prior knowledge of the parameters, except for the fact that they exist; thus, $p_1(y/m)$ explicitly incorporates model nonlinearity and knowledge of model uncertainties in simulating the measurements and does not include the knowledge of uncertainty in the parameters.

The probability density $p_2(m, y)$ results from information contained in the joint space of the control parameters and measurements without knowledge of the model that maps the parameters into observation space. This implies that the parameters in principle have existence or definition outside the particular model that is employed in the inverse problem. For example, a rate of surface emission of trace gases exists regardless of whether a small-scale turbulence or a global circulation model requires its specification. To the extent that the parameters are geophysical quantities or derivatives thereof, there will exist a relationship

between parameters and observations dictated by the physics of the system of interest. Clearly, if this relationship does not exist, no meaningful solution to the inverse problem is possible.

If we assume that the model consists of a description of a physical system, then prior to solving the inverse problem using a given set of observations, model parameters cannot include information about this set of observations. It is, therefore, natural to assume that the information about the parameters and observations is independent from each other. This implies that

$$p_2(m, y) = p_2(m)p_2(y). \quad (2.5)$$

The probability density $p_2(y)$ arises exclusively from information about the uncertainties or errors in the measurements; the probability density $p_2(m)$ consists of a priori information on the parameters.

Under the same assumption used in (2.5), the joint homogenous probability density reads

$$v(m, y) = v(m)v(y). \quad (2.6)$$

Substituting (2.4)–(2.6) into (2.2) leads to

$$p(m, y) = \frac{1}{\gamma} \frac{p_2(m)p_2(y)p_1(y/m)}{v(y)}, \quad (2.7)$$

and using (2.7) in (2.3) renders the final solution to the inverse problem as the posterior parameter probability density distribution

$$p_m(m) = \int_D \frac{1}{\gamma^*} [p_2(m)p_2(y)p_1(y/m)] dy. \quad (2.8)$$

The homogenous PDF $v(y)$ is by definition the probability associated with a unit discrete volume in the space of all possible measured values. This probability is constant over the range of y as long as the measurement space is linear, so $v(y)$ is folded into new constant γ^* for convenience.

The relationship (2.8) shows that the contributions from prior, measurements, model nonlinearity, and the associated modeling errors to the posterior marginal PDF are uniquely identifiable by specifying each of the PDFs on the right-hand side. In the next section, we show the theoretical result of assuming that each contributing PDF is Gaussian; however, in the numerical analysis method in section 3 and in the examples in section 4, these PDFs are explicitly numerically evaluated and are not necessarily assumed to be Gaussian.

b. Gaussian prior, measurement, and model probability densities

As was mentioned in the introduction, it is common to assume that the probability density functions associ-

ated with modeled measurements, actual measurements, and prior information about the parameters are Gaussian. In the joint measurement space D these can be written

$$p_1(y/m) = \frac{1}{(2\pi) \det|\mathbf{C}_s|^{1/2}} \times \exp\left\{-\frac{1}{2}[y - f(m)]^T \mathbf{C}_s^{-1}[y - f(m)]\right\} \quad \text{and} \quad (2.9)$$

$$p_2(y) = \frac{1}{(2\pi) \det|\mathbf{C}_y|^{1/2}} \times \exp\left[-\frac{1}{2}(y - \bar{y})^T \mathbf{C}_y^{-1}(y - \bar{y})\right], \quad (2.10)$$

where $\mathbf{C}_s$ and $\mathbf{C}_y$ are the model and measurement covariance matrices, respectively. The Gaussian form defined in (2.9) represents the PDF of simulated observations for every discrete m in the space Π . This property implies that $p_1(y/m)$ is at least a two-dimensional function in the space $(\Pi \times D)$ as shown in section 4. It is important to note that the stochastic Gaussian quantity in (2.9) is the simulated measurement with the expectation value (i.e., the mean) equal to the model solution for each m . The stochastic nature of y in (2.9) is thus the consequence of model uncertainty, which results not from uncertainty in the control parameters but from other sources of model error. The examples in section 4 clearly illustrate this feature, which is critical for understanding the contribution of the model nonlinearity and the modeling errors to the posterior PDF. Actual errors in the earth system models are likely not Gaussian, but this relatively simple form of model error serves to illustrate the role of model error in the nonlinear inverse problems.

The stochastic Gaussian quantity in (2.10) is the actual physical measurement with mean value $\bar{y}$ denoting the statistical mean of the physical measurements. Here there is no reference to the true value of the measurement because the truth is not known. Consequently, the value $\bar{y}$ is to be understood simply as an estimate of the statistical mean of the measurement, without reference to whether or not it has bias relative to the truth. For each measurement type, the convolution of (2.9) and (2.10) (i.e., the multiplication of PDFs in the same space) and insertion in (2.8) generates (Tarantola 2005)

$$p_m(m) \approx p_2(m) \exp\left\{-\frac{1}{2}[f(m) - \bar{y}]^T \mathbf{C}_D^{-1}[f(m) - \bar{y}]\right\}, \quad (2.11)$$

where the covariance that results from the convolution is

$$\mathbf{C}_D = \mathbf{C}_y + \mathbf{C}_s. \quad (2.12)$$

Expression (2.11) indicates that the convolution of two Gaussian PDFs in the measurement space is also Gaussian, with summed uncertainties derived from independent modeled and measured information and represented by the covariance $\mathbf{C}_D$.

Assuming further that the prior PDF in the parameter space is also Gaussian,

$$p_2(m) = \frac{1}{(2\pi) \det|\mathbf{C}_m|^{1/2}} \times \exp\left[-\frac{1}{2}(m - m_{\text{prior}})^T \mathbf{C}_m^{-1}(m - m_{\text{prior}})\right], \quad (2.13)$$

it follows that

$$p_m(m) \approx \exp[-J(m)], \quad \text{and} \quad J(m) = \frac{1}{2}[[f(m) - \bar{y}]^T \mathbf{C}_D^{-1}[f(m) - \bar{y}] + (m - m_{\text{prior}})^T \mathbf{C}_m^{-1}(m - m_{\text{prior}})]. \quad (2.14)$$

It is well known that unless $f(m)$ is linear, the posterior $p_m(m)$ is not a Gaussian PDF. If f is linear and denoted $\mathbf{H}$, it is possible to show that the posterior PDF is Gaussian with the following first and second moments (Tarantola 2005) (the mean and covariance, respectively):

$$\begin{aligned} \langle m \rangle &= m_{\text{prior}} + \mathbf{C}_m \mathbf{H}^T (\mathbf{H} \mathbf{C}_m \mathbf{H}^T + \mathbf{C}_D)^{-1} \\ &\times (\bar{y} - \mathbf{H} m_{\text{prior}}), \quad \text{and} \\ \mathbf{C} &= (\mathbf{H}^T \mathbf{C}_D^{-1} \mathbf{H} + \mathbf{C}_m^{-1})^{-1}. \end{aligned} \quad (2.15)$$

The equivalent of expression (2.15) is used in the formulation of the set of ensemble data assimilation techniques that derive from Kalman filter theory, where m is referred to as the state of, for example, the atmosphere or ocean, at an instant in time τ when the measurements are available; m_{prior} is the mean of the state prior to using the measurements at time τ ; and $\mathbf{C}_m$ is the associated error covariance. These are obtained from a forecast model by propagation of initial conditions from time $\tau - \Delta\tau$, where $\Delta\tau$ is the interval between observations. The linear model $\mathbf{H}$ is typically referred to as the observational operator because it is used to transform the prior state into the observation space. In KF-type techniques, the mean $\langle m \rangle$ in (2.15) is referred to as the minimum variance estimate of the state $m(\tau)$, and $\mathbf{C}$ is the new error covariance of that

state given the observation, model, and prior uncertainties (Evensen 2006; Lewis et al. 2006; Kalnay 2003). In the KF approach, model uncertainties are accounted for in two terms: the errors in the observational operator model are included in $\mathbf{C}_D$ and errors associated with the forecast model are included in the prior $\mathbf{C}_m$.

It is well known that when the model is linear, (2.15) is also the solution of $\arg \min[J(m)]$. In spatial variational data assimilation techniques (e.g., 3DVAR)—which, as in the KF-based techniques, solve for the state $m(\tau)$ —an expression analogous to (2.15) is used to indicate the form of the final solution of the data assimilation problem, but the solution is obtained via minimization of the cost function $J(m)$. Finally, data assimilation techniques that minimize a cost function with respect to the initial condition instead of with respect to $m(\tau)$, such as 4DVAR techniques, define m in (2.14) and (2.15) as the initial control state. In the formulation of 4DVAR techniques, the equivalent of expression (2.15) is also used to indicate the form of the final solution; in this case, the linear model $\mathbf{H}$ includes the combined effect of both the forecast and observational operator models. This implies that in the 4DVAR data assimilation formulation the modeling errors of both the forecast and observation model are included in $\mathbf{C}_s$, as in (2.9).

When the model in (2.14) is nonlinear, which is more typical of geophysical inverse problems, the KF-based technique implicitly assumes that the nonlinear mapping into the observation space by the observational operator renders an approximately Gaussian posterior PDF. Because the forecast model in KF-type techniques is involved in the prior, the assumption of a Gaussian prior PDF implies that propagation of the Gaussian PDF in the initial condition by the nonlinear forecast model produces an approximately Gaussian prior at time τ . 3DVAR-type techniques include the same assumption for the prior as in KF-type techniques, but the posterior is not required to be Gaussian when minimizing the cost function, which by definition includes nonlinear mapping into the measurement space. The posterior is, however, often implicitly assumed approximately Gaussian in the 3DVAR techniques when it is used as the initial condition for a forecast that results in a new prior for the subsequent data assimilation step.

In the 4DVAR techniques, the nonlinear $f(m)$ in (2.14) is the mapping of the initial condition into measurement space via the combined effect of the forecast model and observational operator. The inverse solution is obtained by minimization of the cost function, and in this way (in theory) the result of data assimilation is the maximum of the posterior PDF, which is not required

to be Gaussian or any other particular form. This feature of the 4DVAR-type technique is similar to 3DVAR techniques, except that the nonlinear mapping includes the forecast model in addition to the observational operator. The critical underlying belief in variational techniques (if not an explicit assumption) is that the posterior PDF is at least unimodal, if not Gaussian, within the considered range of parameter values because the posterior PDF would then have only one maximum, which would guarantee a well-posed minimization problem. Similar to both the KF- and 3DVAR-type techniques, when the state of the system must be updated sequentially, as for example in atmospheric state analysis, the 4DVAR type technique implicitly includes the assumption that the prior PDF is Gaussian. Unlike the KF- and 3DVAR-type techniques, however, the 4DVAR technique does not update the prior with the Gaussian approximation at every observation time because sequences of observations over a prescribed interval are assimilated at the same time. The length of the assimilation interval is determined by practical constraints but could, in principle, be long. Contrary to what is commonly believed, the analysis in this study shows that the assimilation interval is not determined by a requirement that the forecast model should be approximately linear within the assimilation interval.

In the numerical method which is described in the next section, the joint posterior PDF given by (2.2) and its marginal posterior (2.8) are evaluated by explicit computation without imposing any of the approximations used in modern data assimilation techniques. In this way, the full effect of model nonlinearity and the modeling and observation uncertainties on the posterior PDF can be investigated, from which it is then possible to determine the validity of standard approximations.

3. Numerical method for explicit evaluation of the posterior PDF

In the expression for the posterior PDF (2.8), Monte Carlo sampling can be used to numerically evaluate each PDF on the right-hand side if the transformation from the parameter into observation space is known and the stochastic quantities are each given a prescribed distribution. Let us start with the model PDF $p_1(y/m)$. The transformation from m into y is given by the model $f(m)$, and it is possible to define a discrete grid in the space m composed of points with coordinates in each phase space direction (i.e., for each physical quantity in m). A point on this grid is written

$$m_i = (m_i^1, \dots, m_i^M), \quad (3.1)$$

the grid coordinate in each phase direction j is

$$m_i^j = m_{\min}^j + (i - 1)\Delta m^j, \quad (3.2)$$

and grid coordinates are defined over finite intervals $[m_{\min}^j, m_{\max}^j]$. The unit grid interval in each direction is $\Delta m^j = \text{const}$, and we note that this constant is typically different for each phase direction because each direction often represents a different physical quantity. The simulated measurements are then determined on this regular grid by

$$y_i = f(m_i). \quad (3.3)$$

By this transformation the nonlinear functional dependence of the simulated measurements on the parameters is mapped. This is equivalent to producing a graph of a function on a coordinate grid.

The next step in constructing $p_1(y/m)$ is to assign a PDF to each of the simulated measurements for every m point. To make the analysis tractable, it is assumed that the PDF that represents the model errors in f is Gaussian, as in (2.9). This assumption is not necessarily very restrictive because any amount of variance could be used to account for the model discrepancy at each m , including structural errors in the model relative to the real physical system from which the actual measurements are taken. Assigning a Gaussian PDF to each m on the grid is equivalent to assigning a normally distributed random sample in the y space for each y_i , with the mean equal to y_i and with a prescribed variance. In this study, standard Monte Carlo sampling from the normal distribution with prescribed mean and variance is applied. The use of Monte Carlo sampling implies a very large sample for each y_i ; on the order of 10^4 . The penultimate step in constructing $p_1(y/m)$ in the joint space $D \times \Pi$ is to define a discrete and regular grid in the measurement space using

$$y_l = y_{\min} + (l - 1)\Delta y \quad (3.4)$$

on the finite interval $[y_{\min}, y_{\max}]$. The interval boundaries are chosen such that all simulated and actual measurement values are included. The final step in evaluating $p_1(y/m)$ is to place the Monte Carlo samples from the Gaussian distribution on the regular y grid. This renders the final discrete map of $p_1(y/m)$ in the joint space $D \times \Pi$. The examples of $p_1(y/m)$ for different f in section 4 clearly illustrate the result of this procedure. As expected, the results are easiest to visualize in two dimensions, with only one physical quantity in the measurements and one in the parameters.

The next PDF in (2.8) to be numerically evaluated is the one that corresponds to the actual measurements, $p_2(y)$. Because this PDF is combined with $p_1(y/m)$ in (2.8) (i.e., the two PDFs are multiplied in the same

space) it must be discretized on the same regular grid (3.4). This is done by first generating a large Monte Carlo sample from the prescribed distribution, placing the samples onto the regular grid, and computing the PDF by normalizing the count in each unit interval. It should be noted that the prescribed distribution for the actual measurements could be assumed to take the form of any known statistical distribution as long as the distribution parameters are provided. In the examples in section 4, the Gaussian PDF is used with varying values of mean and variance for each measurement type. This is appropriate both because the examples in this first part of the study are rather tutorial in nature and because the physical quantities used for the measurements allow for the use of the normal distribution.

The prior PDF in (2.8) demands special attention. To begin with, it is first assumed that the prior PDF is homogenous in each direction in the m space. This is essentially equivalent to assuming that each parameter can take on any value within a physically realistic range with equal probability. Because the uniform distribution is not as simple to manipulate analytically, samples for the homogenous prior distribution are generated via Monte Carlo sampling using a Gaussian distribution with very large variance, and with a mean equal to an arbitrary value of m_i^j drawn from within the grid bounds in each direction independently. This sample is then placed on the regular grid defined in (3.2), and the multidimensional PDF is evaluated by normalization of the count in each interval on the grid. Because the starting prior PDF is homogenous, the joint prior for all directions combined can be written

$$p_2^0(m) \equiv p_2^0(m^1, \dots, m^M) = p_2^0(m^1)p_2^0(m^2) \dots p_2^0(m^M), \quad (3.5)$$

where the superscript 0 indicates the starting prior. The right-hand side of (3.5) formally indicates independence of the marginals in each direction. The prior (3.5) is combined by conjunction with the product $p_2(y)p_1(y/m)$ in the (y, m) space via multiplication of PDFs on a regular grid in each phase dimension of m for every y grid interval.

This procedure produces the posterior PDF that results from a single measurement instance, which is written

$$p(y^1, m) = \frac{1}{\gamma^*} [p_2^0(m)p_2(y^1)p_1(y^1/m)]. \quad (3.6)$$

When subsequent measurements are added, the prior PDF is updated with the previous posterior PDF which, after the first instance, is no longer uniform but is instead the full joint multidimensional PDF. The

marginal posterior in m space is then computed by evaluating

$$p_m(m) = \sum_i p(y_i^1, m) \Delta y,$$

and it follows that

$$p_2^1(m) \equiv p_m(m).$$

At every measurement after the first one, the posterior PDF is computed by

$$p(y^n, m) = \frac{1}{\gamma^*} [p_2^{n-1}(m) p_2(y^n) p_1(y^n/m)]$$

$$p_m^n(m) = \sum_i p(y_i^n, m) \Delta y, \quad (3.7)$$

where superscript n denotes the n th measurement instance. The observation instances in the examples in section 4 are all instances in time because of the simplicity of the models used. Observation instances in space and/or in different measured physical quantities would be treated the same way when updating the prior; however, in either distribution of observations, the underlying assumption in (3.7) is that the measurements are mutually independent. To include dependency between observations, it would be necessary to use, for example, the covariance C_y when constructing the Gaussian PDF in the actual measurement space. This could be done by convolution of $p_2(y)$ with $p_1(y/m)$ in all y directions at the same time. This is not strictly necessary for this first analysis of the impact of model nonlinearity and is thus deemed to be beyond the scope of the current study. In addition, we note that the assumption of statistical independence between measurements is common among data assimilation techniques used in the earth system sciences.

To test the effect of approximating the prior with a Gaussian instead of the actual PDF, some of the experiments described in section 4 employ a joint multi-dimensional Gaussian prior PDF computed from the actual marginal posterior in (3.7). This is done by first evaluating the mean and covariance of the joint PDF in parameter space $p_m^n(m)$ in (3.7) and then using these to compute the joint Gaussian PDF in the same space. The Gaussian prior approximation is used in section 4 in each experiment that is intended to test features of either the KF or 3DVAR type of technique. We wish to emphasize that when numerical experiments are used to test the properties of the class of data assimilation techniques examined in the next section, this is done in a theoretical manner. That is, experiments are performed for the purpose of testing a particular theoretical approximation, and practical limitations necessary

for implementation of modern techniques are not applied. The numerical computations in the current study all rely on large Monte Carlo samples and do not assume a small sample size [as in the ensemble Kalman filter (EnKF) and ensemble Kalman smoother (EnKS) algorithms], nor do they employ a minimization technique or computation of cost function gradients as would be done in a variational technique.

4. Examples

a. Damped harmonic oscillations

To demonstrate the basic properties of the numerical analysis method described above, we begin with a simple model of damped oscillations. The governing equation can be written

$$\frac{d^2\chi}{d\tau^2} + \alpha \frac{d\chi}{d\tau} + \omega\chi = 0, \quad (4.1)$$

where χ is the displacement from an equilibrium point (also called the oscillation amplitude), τ is time, α is a damping coefficient which represents a dissipative force (e.g., friction or air resistance), and ω is the natural frequency coefficient (e.g., buoyancy). This model (4.1) has been used in many branches of physics to simulate the evolution of a particle or mass moving through a region whose potential has one or more local minima, including motion of a pendulum, planetary and satellite motions, the classical description of an electron in orbit around a nucleus, and an air parcel in a geopotential field, to mention a few.

The solution to (4.1) is readily obtained as

$$\chi = A_1 e^{(-\lambda + \eta)\tau} + A_2 e^{(-\lambda - \eta)\tau}, \quad (4.2)$$

where A_1 and A_2 are determined from initial conditions and

$$\lambda \equiv \frac{\alpha}{2},$$

$$\eta = \sqrt{\lambda^2 - \omega_0^2}, \quad \text{and}$$

$$\omega_0^2 \equiv \omega.$$

Using measurements of oscillation amplitude, there are two types of parameters for which the inverse problem in (4.2) could be defined: initial conditions and coefficients. Each problem conveniently demonstrates the properties of the analysis method described in section 3 because each includes two easily recognized transfer functions. In the case of initial conditions, the measurements have a linear dependence on the control parameters; in the inverse problem for the coefficients, the function is nonlinear and exponential. In each problem,

FIG. 1. Probability density functions for the initial condition inverse problem using the damped oscillations model. (a) PDF of observations, shown as the unnormalized count, for oscillation amplitude; (b) joint model PDF for oscillation amplitude on vertical axis and control parameter values on horizontal axis; (c) joint posterior PDF, normalized, with the same axes as in (b), and (d) marginal PDFs, normalized, with initial condition values on the horizontal axis; the solid curve indicates the marginal posterior, the dashed curve indicates the marginal prior, and the solid bar indicates the true value of the initial condition. In (b) and (c) the maximum probability is shown in red and minimum in dark blue; exact values are irrelevant.

amplitude is the measured quantity, and only one observation is taken at an arbitrary time. Figures 1a and 2a show the assumed Gaussian PDFs of the measurements $[p_2(y), (2.10)]$. In these experiments, the mean value of the measurements was produced using $A_1 = 1$, $A_2 = 0$, $\alpha = 0.2$, $\omega = 3$, and $\tau = 10$ in (4.2), and the standard deviation of the measurements is set to $\sigma_y = 0.2$.

The joint PDFs of the modeled measurement and the control parameter $[p_1(y/m)]$ are depicted in Figs. 1b and 2b for the initial condition A_1 and the coefficient α , respectively, for which it is assumed that the model errors are Gaussian for each discrete parameter value. The measurement and parameter spaces are discretized using 50 evenly spaced bins within intervals depicted in Figs. 1b and 2b. We assume the model error is unbiased Gaussian for the purpose of focusing the analysis on the impact of the model function, as discussed in section 2b. The results depicted in Figs. 1b and 2b clearly indicate that the joint model PDF takes the shape of the transfer function. When this PDF is convolved with the measurement PDF and combined (via conjunction) with the prior—which is assumed wide Gaussian and nearly uniform (dashed curves in Figs. 1d and 2d)—the resulting joint PDF in the joint space (y, m) is two-dimensional

Gaussian for the linear problem (Fig. 1c) and two-dimensional joint Gaussian and approximately lognormal for the exponential problem (Fig. 2c). The marginal posterior PDF in the parameter space (solid curves in Figs. 1d and 2d) is Gaussian for the initial condition problem (Fig. 1d) and approximately lognormal for the coefficient problem (Fig. 2d), a result that arises naturally from the form of the model function (4.2), the inverse relationship $\alpha(\ln x)$, and the fact that measurements of oscillation amplitude are normally distributed in the bounded domain. To more clearly illustrate the fact that the PDF of α is approximately lognormal, we display the posterior marginal PDF in $\ln(\alpha)$ space in Fig. 2d; the transformed PDF is nearly Gaussian centered on the true solution. The specified true value of the control parameter, depicted by vertical bars in Figs. 1d and 2d, is located at the mean of the marginal posterior in the case of the linear transfer function and at the mode for the nonlinear exponential transfer function.

These examples of solutions to an inverse problem using a simple model with a known analytical solution demonstrate that the theory presented in MT02 and the corresponding numerical method described in section 3 together enable an explicit account of the contribution

FIG. 2. Same as Fig. 1, but for the damping coefficient (α) inverse problem. In (d) the PDF of $\ln(\alpha)$ is shown to demonstrate that the posterior PDF for α is nearly lognormal.

of each source of information in the inverse problem and result in the expected posterior solution. In the next section, Lorenz's three-variable model is used to study the effect of model nonlinearity on the posterior PDF for the case of a multidimensional parameter space when the analytical form of the transfer function is unknown.

b. Lorenz three-variable model

The experiments in this section are performed using the well-known Lorenz three-variable model (Lorenz 1963). This model was designed to simulate atmospheric dry cellular convection in a layer of fluid with uniform depth embedded in a gravitational field and with a constant temperature difference between the upper and lower layer boundary. The model simulates the evolution of a forced dissipative hydrodynamic system that possesses nonperiodic and unstable (i.e., chaotic) solutions. The governing equations for this model can be written

$$\begin{aligned} \frac{dX}{d\tau} &= -a(X - Y), \\ \frac{dY}{d\tau} &= rX - Y - XZ, \quad \text{and} \\ \frac{dZ}{d\tau} &= XY - bZ, \end{aligned} \quad (4.3)$$

where X , Y , and Z are the three state components. The control parameters in the model are the initial conditions (X^0 , Y^0 , Z^0) and the physical coefficients (a , r , b).

The model Eq. (4.2) is solved numerically using a Runge–Kutta finite differencing scheme. The reference, or true, values of the physical coefficients are set equal to those used in the original Lorenz (1963) study: $a = 10$, $r = 8/3$, $b = 28$. Several different sets of control initial conditions were tested in the experiments to evaluate the expected independence of the analysis results to the particular period of the state evolution, and it was found that the effects of the nonlinear model on the posterior PDF were not dependent on the choice of initial condition (not shown). The results presented below are examples taken from experiments in which the true solution was obtained using $X_0 = Y_0 = Z_0 = 1$. In each experiment the total simulation time was $40 \Delta\tau$, where $\Delta\tau = 0.04$ is the nondimensional time step. The measurements were derived from the true solution of one of the components of the state (X , Y , Z) at evenly spaced time points. In each of the examples discussed in the following sections, the frequency of the measurements is $10 \Delta\tau$; this measurement time interval was found by experimentation to provide sufficient constraint on the inverse solution for all control parameters for the particular choice of the true coefficients. The stochastic measurements at each observation time were generated by Monte Carlo Gaussian sampling with the

mean equal to the true solution and with standard deviation of two.

Because the measurements are equivalent to a component of the state, the observational operator defined in section 2b is defined as (e.g., for component X)

$$\mathbf{H} \equiv \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

This choice was made to enable examination of the effect of a dynamical nonlinear model (4.3) without introducing complications associated with an additional transformation from state into observation space. In applications of data assimilation techniques in the earth system sciences, the observational operators are typically much more complex and often nonlinear, but the effect of this nonlinearity is expected to be qualitatively similar to the equivalent nonlinearities in the dynamical model. In addition, the use of the simplest possible observational operator enables straightforward testing of the assumption of the Gaussian prior, which facilitates a more robust evaluation of the properties of the state estimation problem with respect to the model uncertainties.

Numerical experiments with the Lorenz's model are performed for the following inverse problems: (i) single coefficient or initial condition, (ii) all three coefficients, (iii) all three initial conditions, and (iv) state (X, Y, Z) at the measurement times. In each inverse problem, identical measurement times are used, and model and measurement uncertainties are varied from small to large to test their impact on the effect of the nonlinear model on the posterior PDF. The effect of the Gaussian PDF assumption for the prior was also tested in each inverse problem. The numerical experiments are summarized in Table 1.

1) NONLINEARITY AND OBSERVATIONAL CONSTRAINT

The effect of the nonlinear forward model, which provides the nonlinear transfer function between the control parameters and observations, is easiest to analyze by first considering the inverse problem for single parameters (experiments 1 and 2 in Table 1). This type of result has already been discussed in section 4c for the example of the damped oscillations model, in which natural frequency was used as the control parameter. The analysis with the Lorenz model in this section provides an extension to a model with much more complex nonlinearity, the character of which changes over time. The time variability of the transfer function allows the opportunity to evaluate the impact of observations

TABLE 1. Numerical experiments with Lorenz's model.

Experiment	Control parameter	Prior PDF	Model error
1	a	Uniform	Moderate
2	X_0	Uniform	Moderate
3	a	Exact	Moderate
4	X_0	Exact	Moderate
5	(a, r, b)	Exact	Negligible
6	(a, r, b)	Exact	Moderate
7	(a, r, b)	Exact	Large
8	(X_0, Y_0, Z_0)	Exact	Negligible
9	(X_0, Y_0, Z_0)	Exact	Moderate
10	(X_0, Y_0, Z_0)	Exact	Large
11	(a, r, b)	Gaussian	Moderate
12	(X_0, Y_0, Z_0)	Gaussian	Moderate
13	(X^n, Y^n, Z^n)	Exact	Moderate
14	(X^n, Y^n, Z^n)	Gaussian	Negligible

from different instances in time, individually and cumulatively by the prior update, as described in section 3. The probability density functions associated with the observations, prior, model, and posterior for each observation time are shown in Figs. 3 and 4 for coefficient a and initial condition X_0 , respectively. It is important to note that in this first set of experiments, information from different times was considered in isolation; the effect of accumulating information at sequential times is considered below.

Results from the inverse problem for coefficient a and for each observation instance individually are depicted in Fig. 3. PDFs of the observations were specified Gaussian, assumed to be uncorrelated in time, and assigned a constant variance (Figs. 3a,e,i,m). The joint PDF that results from mapping the parameters on a regular grid into observation space (Figs. 3b,f,j,n) exhibits increasing complexity with time, demonstrating the accumulated effect of nonlinearity in the forward model solution. Convolution of the model PDF with the Gaussian observation PDF at each observation time and conjunction of the result with the noninformative prior, results in a bimodal joint posterior PDF (Figs. 3c,g,k,o). The associated marginal PDFs (Figs. 3d,h,l,p) are also multimodal and exhibit increasing complexity with time. For reference, each panel with the marginal posterior PDF includes the delta function PDF of the exact true value of the coefficient, shown as a single vertical bar.

Several important properties of the nonlinear inverse solution can be seen in the results. First, multimodality in the posterior PDF is directly related to nonmonotonicity in the model solution, not just to nonlinearity in the model. Second, the distinct modes in the posterior PDF reside in a joint space that consists of regions in the parameter space with steep gradients $\partial f/\partial a$ and re-

FIG. 3. PDFs for the a coefficient inverse problem without prior update (Lorenz's model; expt. 1 in Table 1) at observation times 1, 2, 3, and 4, top to bottom, in each column. The first column displays the unnormalized observation PDFs; the second column, the joint normalized model PDFs; the third column, the joint normalized posterior PDFs; and the fourth column, the normalized marginal PDFs. Curves and solid bar as in Fig. 1; exact values are irrelevant.

gions in the observation space with nonzero values of the observation PDF (depicted as blue horizontal lines in the model PDF panels in Figs. 3 and 4). It can also be seen that the strength of the gradient in the parameter space is directly related to amplitude of the PDF in the associated mode (e.g., cf. Figs. 3d with 3p). This property implies that the parameter values to which the model function is most sensitive will tend to dominate the solution locally in the parameter space, which could lead to the incorrect maximum likelihood estimate in practice if the global maximum is desired. This circumstance is equivalent to the potential for converging to a local minimum of the cost function (2.14). It is important to notice that the potential for an incorrect global maximum likelihood solution is a consequence of f being nonmonotonic nonlinear rather than just nonlin-

ear, as commonly implied in the data assimilation literature.

The third important conclusion is that there are multiple regions in the parameter space which have similar probability at all times; very different parameter values produced similar values of the forward observations. This property implies that each observation individually did not provide enough information to identify the true mode. This is a direct consequence of the fact that the problem is nonmonotonic nonlinear and not just nonlinear, in contrast to the damped oscillator model examined in section 4a. The different shape and modality of the posterior PDF at different observation instances with the nonmonotonic nonlinear model implies different observational constraint from equally accurate observations, where observational constraint is

FIG. 4. Same as Fig. 3, but for the X_0 initial condition inverse problem (Lorenz's model, expt. 2 in Table 1).

defined as information contained in the observation that contributes toward identifying the correct solution. This information is constant for different observation instances with the linear and monotonic nonlinear f in section 4a. The results in Figs. 3d and 4d also illustrate another important property of the observational constraint in the case of a nonmonotonic nonlinear forward model: the number of different observation instances needed to identify the correct solution is larger than the number of degrees of freedom in the control parameter space. In the current simple examples, the number of degrees of freedom is 1 (i.e., the coefficient a or initial condition X_0). This result implies that a greater number of independent observations is needed to constrain the solution for a nonmonotonic model than are required for a monotonic model for a given number of degrees of freedom.

In the experiments with the initial condition for the Lorenz model (experiment 2 in Table 1) joint model

PDFs are less complex than for the coefficient a and primarily exhibit two distinct extremes (Figs. 4b,f,j,n). These extremes arise from the two modes of the well-known chaotic attractor for the Lorenz 1963 model. Despite less complexity in the joint model PDFs, convolution of the result with the observation PDFs produces a variety of different posterior PDFs, the shapes of which depend on the location of the observations in the phase space at each time. For example, the posterior marginal PDF was trimodal at the first observation time (Fig. 4d), virtually uniform at the second and third times (Figs. 4h,l), and bimodal at the last time (Fig. 4p). The corresponding results for other two initial conditions were similar (not shown), and we note that the same general relationship between posterior PDFs and the properties of the nonlinear model function (i.e., whether or not the relationship is monotonic and the strength of the gradient in parameter space) can be seen in the initial condition results, as for the coefficient a .

FIG. 5. PDFs for the a coefficient inverse problem with prior update (Lorenz's model, expt. 3 in Table 1) at observation times 3, 4, and 5, top to bottom, respectively, in each column. The first column displays the joint posterior PDFs with the parameter values on the horizontal axis and the observation values on the vertical axis and the second column displays the marginal PDFs; the solid curve indicates the posterior marginal; the dashed curve, the prior marginal; and solid bar, the true solution.

The next step in the analysis was to sequentially combine information from different observation instances, as would be normally done in an inverse problem with multiple data points or in data assimilation. The assumed mutually independent observations are included in the posterior solution via the procedure described in section 3, specifically through application of Eq. (3.7). That is, at each observation instance after the first, the prior PDF is updated using the previous posterior. It is important to note that the model PDF is not changed by this procedure because it is defined as the joint PDF in the combined measurement and parameter space, in which the parameters are characterized by the homogeneous marginal PDF (i.e., the uniform PDF) before the prior and observations are taken into account. Because the observation and joint model PDFs are identical to

those in Fig. 3, we will only discuss the marginal and joint posterior solutions.

The effect of updating the prior with each added observation is illustrated in Fig. 5 for the inversion of coefficient a (experiment 3 in Table 1). The results show that the posterior PDFs, both joint and marginal, become more localized with each update, resulting in virtually unimodal but nonsymmetric PDFs (Figs. 5e,f). It is interesting to note that the solution appears qualitatively similar to that in Fig. 3 at the last two observation times, with the notable exception of the absence of the secondary mode in the solution. It appears that accumulation of observation information in time has rendered the problem better posed; there is now a unique modal solution to the inverse problem. As such, the final posterior PDF in Fig. 5f provides the definition

FIG. 6. Marginal posterior PDFs for the coefficients of the Lorenz model from the experiments with the exact prior update and moderate and large model errors (solid and dashed curves, respectively) and from the experiment with the Gaussian prior update and moderate model errors (dotted curves) PDFs for (a) coefficient a , (b) coefficient r , and (c) coefficient b . The gray bar shows the bin location that contains the true solution for each coefficient.

of the data assimilation problem from the point of view of the techniques discussed in section 2. The problem in this particular example is to solve for the properties of the unimodal, nonsymmetric posterior PDF. In contrast to the solution for the coefficient a , the solution for the initial condition X_0 (experiment 4 in Table 1) demonstrated multiple modes at early times but rapidly converged to a narrow and unimodal PDF centered around the true solution (not shown). The implication is that when the X component of the state is observed, the inverse solution for the initial condition in X is nearly perfectly constrained, despite nonmonotonic nonlinearity and model error. In this case, the data assimilation problem would be trivial to solve for the given set of observations.

2) IMPACT OF MODEL ERROR

In the examples presented in the previous sections, model uncertainty was specified to be Gaussian with a constant variance. In this section, the impact of the model error amplitude on the posterior PDF is tested by varying the variance in the discrete representation of (2.9). Experiments were performed for the single and multiparameter inverse problems, but only results from the multiparameter experiment are shown (experiments 5–12 in Table 1). All experiments included an update of the prior with each added observation, and the variance for each model error $y^n(m)$ was set to 0.000 01, 1 and 16, intended to represent negligible, moderate, and large error, respectively [$y^n(m)$ is defined in

section 3]. It should be noted that the moderate model error variance was used in the single coefficient and initial condition experiments and that the resulting level of uncertainty can be seen in the corresponding joint model PDFs (Figs. 3d,h,l,p and 4d,h,l,p). The discretization in phase space of the parameters in the multiparameter experiments is 10 evenly spaced bins in each direction.

Marginal posterior PDFs for the model coefficients (a , r , b) are depicted in Fig. 6 for moderate and large model error amplitudes, respectively. For the experiment with negligible model error (experiment 5) the final-time posterior marginal PDFs were almost identical to the perfect-knowledge PDFs of the true parameter values and exhibited a deltalike function at the true location (not shown). Marginal posterior PDFs for coefficients b and r were characterized by multimodality at early observation times but converged to the delta function after observations from all four times were included.

Increasing the model error to moderate values for the multidimensional coefficient problem (experiment 6) has the effect of widening the posterior marginal PDFs for all coefficients (solid curves in Fig. 6). In addition, the final marginal PDF for the coefficient r (solid curve in Fig. 6b) exhibits multimodality, with two distinct modes and higher cumulative probability at values lower than the true solution, though the mode with slightly higher probability coincides with the true solution. Application of large model error (experiment 7)

FIG. 7. Same as Fig. 6, but for the initial conditions from the experiments with the Lorenz model PDF for the (a) X , (b) Y , and (c) Z component of the initial condition. The gray bar shows the bin location that contains the true solution for each component.

further increases the width of the marginal PDFs for all coefficients (dashed curves in Fig. 6) and has the additional effect of significantly damping the multimodal character of the solution for coefficient r (dashed curve in Fig. 6b). In the case of the a coefficient, the solution is now decidedly multimodal, and the primary mode is located relatively far from the true solution (dashed curve in Fig. 6a). The results with the large model error suggest that when the model error is assumed Gaussian with large variance to account for potential structural errors in the model, an inverse problem that contains the potential for multimodality (i.e., the nonmonotonic nonlinear model) may exhibit a posterior PDF that has relatively high probability in the region of one of the possible modes—but not necessarily the true mode. In summary, the fact that each coefficient exhibited different behavior with increases in model error serves to illustrate a different aspect of the effect of model error. In the case of coefficient b , the nature of the primary mode in the posterior PDF did not change; instead, an increase in model error simply increased the variance in the posterior. This indicates that although the model is nonmonotonic nonlinear, the solution is approximately monotonic in b ; the nonlinearity appears in the positive skewness evident at all error levels. In the case of coefficient r , each level of error leads to a solution with multiple discrete modes at early observation times (not shown), and the primary mode in the final solution is not easily distinguished. In the case of coefficient a , with small model error only a single mode appears in the posterior PDF, but as the error is increased a sec-

ondary mode becomes evident, and at large values of model error, this mode is the dominant one.

Experiments 8–10, in which model error was varied for the multidimensional initial condition problem, support the same general conclusion but provide additional information about the properties specific to that problem. For example, when moderate model error was added, the marginal PDFs become markedly bimodal for all initial conditions (solid curves in Fig. 7), and a further increase in the model error resulted in wider PDFs with dominance of a secondary mode for the initial conditions Y_0 and Z_0 that is significantly shifted away from the true solution (dashed curves in Figs. 7b and 7c, respectively). These results indicate that, similar to the single initial condition inverse problem, observations of a single component of the state over four time instances are sufficient to constrain the inverse problem for all three initial conditions, but only under the condition of negligible model error.

In summary, the results for multicoefficient and initial condition inverse problems suggest that, in general, the nonmonotonic nonlinear inverse problem could be difficult to solve—not because of nonlinearity in the model, but because of the presence of model uncertainty that falsely increases probability in regions that are separated from the true mode in parameter space but that produces similar values in observation space.

3) IMPACT OF THE GAUSSIAN PRIOR UPDATE

In this section we discuss the impact of assuming that the prior PDF is Gaussian at each observation time. In

experiments 11 and 12 (Table 1), the Gaussian prior is derived from the actual non-Gaussian multidimensional posterior PDF that resulted from the previous observation input, and the model error is assumed to be of moderate amplitude, as in experiments 6 and 9. The a coefficient exhibits a solution that differs greatly from experiment 9, which used an exact prior update. The posterior marginal PDF now contains multiple modes that persist to the final time (dotted curve in Fig. 6a). This behavior results from the fact that the imposed Gaussian prior PDF captures locations of multiple modes by assigning large variance at first observation time to accommodate the width of the exact PDF. The large variance in the prior allows the secondary mode at low coefficient values to persist, similar to the effect of large model error (cf. dashed and dotted curves in Fig. 6a). In the case of the r coefficient, the modes in the exact prior are closely clustered around the truth (solid curve in Fig. 6b), and the assumption of a Gaussian prior has the effect of reducing the solution to a single mode (dotted curve in Fig. 6b) that is located within the true value bin. As in the case with the exact prior, coefficient b exhibits a single well-defined mode close to the true solution, but with larger variance compared with the exact prior (cf. solid and dotted curves in Fig. 6c).

These results reveal the fact that assuming a Gaussian prior can have a wide range of effects on the posterior solution. In particular, for skewed non-Gaussian PDFs with a single mode, the result is a slight bias in the solution and an overestimate of the solution variance (Fig. 6c). For PDFs with a closely spaced secondary mode to one side of the true solution, the Gaussian prior assumption results in a single mode solution (Fig. 6b), whereas a widely spaced secondary mode produces a consistently bimodal solution (Fig. 6a). Stated another way, for a model that produces widely separated modes in parameter space, the Gaussian prior assumption may cause the inverse problem to be difficult to solve accurately, even after incorporating information from additional observations. The results for the initial condition experiments support these same general conclusions (cf. solid and dotted curves in Fig. 7).

4) STATE ESTIMATION

In this section, the impact of nonlinearity, model uncertainty, and assumption of a Gaussian prior is discussed for the problem of state estimation—the inverse problem for the dynamical model state at every observation time. In applications in the atmospheric and oceanic sciences, this type of problem is often addressed using a 3DVAR-type technique or a technique based on KF theory. In the state estimation problem, unlike

the problems addressed in sections 4b(1)–4b(3), model nonlinearity and uncertainty directly affect the prior and indirectly affect the posterior PDF through the conjunction of the observation PDF with the prior PDF. As in previous experiments, the X component of the Lorenz model state was directly observed, and the role of the observational operator $\mathbf{H}$ was not treated. In each experiment, the prior PDF was produced by model integration from the previous observation time. Because the previous observation time serves as the initial condition for the model simulation forward in time, the PDF of the initial condition is set equal to the posterior state PDF at that time. This PDF is used to generate Monte Carlo samples of the initial conditions, which are then integrated forward in time by the model. These samples are subsequently used to compute the prior PDF on a uniform grid in the phase space (X^n, Y^n, Z^n). In the experiments in this section the phase space is discretized in each direction into 15 evenly spaced bins.

The state estimation inverse problem was solved in two ways: (i) the multidimensional prior PDF at each observation time was defined to be the forward model simulation of the exact posterior PDF from the previous time (experiment 13 in Table 1), and (ii) a Gaussian multidimensional PDF was computed from the exact prior (experiment 14). The latter procedure is used in data assimilation techniques based on KF theory, except that in practice samples of the modeled prior stochastic state are typically small relative to the samples used in this study, and the exact PDF is not generated. Instead, first and second moments of the distribution are computed directly from the small sample to characterize the assumed Gaussian posterior PDF. The procedure in (i) is similar to the approach used in particle filter-based data assimilation algorithms (section 1) but differs in two respects: (i) the current analysis explicitly involves discrete representation of the joint observation and parameter spaces, which is discussed in detail in section 2, and (ii) the samples are much larger than the typical sample size used in practical applications of any of the referenced data assimilation algorithms. We wish to point out, again, that the purpose of the numerical experiments in the current study is purely diagnostic. The numerical analysis method used in this study is not to be understood as an inverse problem solution algorithm in and of itself.

As expected, in all experiments the observed component of the state is characterized by a relatively narrow PDF centered at the true solution (solid and dashed curves in Figs. 8a,d,g,j). In contrast, the posterior PDF of the unobserved components of the state varied significantly over the observation times. In experiment 13,

FIG. 8. Marginal PDFs from state estimation experiments with Lorenz's model. PDFs for state components X , Y , and Z are ordered from left to right in each row. The gray bars show the location of the true solution. Observation time increases from top to bottom. Solid curves are for the experiment with exact prior update and moderate model errors (expt. 13 in Table 1); dashed curves are for the experiment with Gaussian prior update and negligible model errors (expt. 14 in Table 1).

with the exact prior update and moderate model error, the posterior PDF for the Y and Z components is wide at the initial time, with a mean in the vicinity of the true solution (solid curves in Figs. 8b,c). At subsequent observation times, the posterior PDF narrows but there is no monotonic progression toward a more accurate solution. For example, at the second time the posterior PDF for Y is unimodal and asymmetric, with a maximum likelihood solution slightly removed from the true solution (solid curve in Fig. 8e), and at the third observation time it is wider, with a maximum likelihood location that is significantly different from the truth (solid curve in Fig. 8h). At the final observation time, the posterior PDF has a modal solution that is located at the true value for all three components of the state (solid curves in Figs. 8j,k,l). These results indicate that the state estimation problem is well constrained at certain observation times and less well constrained at others. Additional experiments with longer observation time periods confirmed this conclusion (not shown), indicating, as may have been expected, that the sequen-

tial state estimation problem does not include the cumulative effect of observations over time, as do the initial condition and coefficient problems.

The change in the posterior PDF over time in experiment 13 was a direct consequence of the change in the nonlinear relationship between state components at each time, together with a change in the associated prior PDF. In sequential state estimation, the prior PDF at each time could be represented as the result of conjunction of two intermediate PDFs: the PDF produced by a model simulation between observation instances and the PDF associated with the nonlinear relationship between observed and unobserved quantities at the end time. This conjunction is written

$$p_2^n(X^n, Y^n, Z^n) = p_2^*(X^n, Y^n, Z^n) p_1^*[X_{\text{obs}}^n / (X^n, Y^n, Z^n)], \quad (4.4)$$

where p_2^* and p_1^* are the first and second intermediate PDFs, respectively, and n is the time index as before. At the third observation time, the marginal posterior

FIG. 9. Joint PDFs for Y and Z state components at the third observation instance for (a) expt. 13 with exact prior update and (b) expt. 14 with Gaussian prior update.

PDF for the Y component does not capture the true solution but instead exhibits a mode at approximately a value of -7 (Fig. 8h) because p_2^* has negligible probability around the true value and p_1^* has a mode at approximately a value of -7 (not shown). The small probability at the true value in p_1^* for Y was a consequence of small probability of the true value in the same component at the previous time (Fig. 8e). These results indicate that even an exact update of the prior PDF in sequential state estimation can lead to loss of probability around the true solution because of non-monotonic nonlinearity in the relationship between state components at the same observation time.

The results obtained from experiment 13 indicate that an accurate maximum likelihood solution can be obtained in the state estimation problem, even in the presence of moderate model error, and suggest that the state estimation problem is, for the given model, less sensitive to model error than the initial condition problem. Although it would be inappropriate to assume a straightforward generalization of this conclusion given the limited experience offered by this relatively simple model, it is encouraging that there is potential for an accurate maximum likelihood state estimation with a nonmonotonic nonlinear model in the presence of model error. It is important to note, however, that there is also the potential for a loss of predictability at the true solution resulting from sequential use of observations that are assumed to be mutually independent.

Finally, the impact of a Gaussian prior update in the state estimation problem was tested in experiment 14 under the condition of negligible model error. The posterior marginal PDFs derived from this experiment are depicted as dashed curves in Fig. 8. The primary impact

of the Gaussian prior update was, as in the experiments in the previous sections, that the posterior PDFs widened and the mean shifted away from the true location for the unobserved state components. The results are similar to those obtained from increasing the model error in the parameter estimation problems in section 4b(3), but in this case the discrepancy between the location of the mean and the true solution remains uncorrected over time. This result indicates that the Gaussian approximation for the prior update does not account for the nonlinear effect of the model, which is represented in the relationship between state components at each observation time. This can be illustrated by comparing the joint posterior PDFs for the state components. Figure 9 shows the joint posterior PDF for the (Y, Z) pair. The joint PDF from the exact prior update experiment in Fig. 9a exhibits an exponential-like relationship between Y and Z along the maximum probability region and a narrow distribution of lower values around this functional relationship. The corresponding Gaussian joint PDF from experiment 13, displayed in Fig. 9b, has the maximum likelihood region in the incorrect location because the relationship between Y and Z is assumed linear, by definition; the Gaussian distribution implies the linear relationship.

5. Summary and conclusions

In this paper, we have described a numerical implementation of the nonlinear inversion technique described by MT02 and have demonstrated its application using two commonly used and well-understood nonlinear models. The underlying theory relies on the conjunction of information from model parameter and ob-

ervation space and assumes that each piece of information can be represented by a probability distribution function (PDF). The numerical technique produces samples of each component PDF for a specific combination of observations and model. In examples using both a damped oscillator model and a model of dry cellular convection (Lorenz 1963), it was shown that the joint PDF of simulated observations generated from a given set of parameter values took the form of the functional relationship between parameters and observations and that the character of the joint posterior PDF of measurements and parameters depended on both the degree of observation uncertainty and on the gradient of the model function in parameter space. Although this result reflects the well-known fact that the characteristics of the posterior PDF depend on the observation uncertainty modulated by the forward model, the numerical technique introduced here allows the possibility of explicitly deconstructing the influence of each component PDF that contributes to the final solution.

After demonstrating the facility of the numerical technique for characterizing the contribution to the posterior marginal and joint distributions by observations and model, the analysis was applied to an examination of model nonlinearity, the role of observations distributed in time, model error magnitude, assumption of a Gaussian prior, and state estimation. The major findings of this study are the following:

- Nonmonotonic transformations from the control parameter into observation space (i.e., the forward model in the inverse problem) give rise to the potential for a multimodal posterior PDF. The realization of this potential depends on the information content of the observations relative to the control state and on observation and model uncertainties.
- The examples with the nonmonotonic Lorenz's model suggest that cumulative effect of observations, over time or space or both, could produce a unimodal final posterior PDF despite strong potential multimodality at each observation instance individually. This result implies that the cumulative effect of observations must be considered to properly evaluate the impact of model nonlinearities in the data assimilation problem.
- The important implication of the potential for a unimodal posterior PD, F through consideration of the cumulative effect of observations in contemporary data assimilation practices, is that the time scale over which the data assimilation problem can be accurately solved for the initial condition maximum likelihood estimate is potentially longer than the time scale within which the linear approximation is valid.

This conclusion implies that linear approximation tests that involve only the forward model are not necessarily useful for evaluating the appropriate time scale over which nonlinear least squares techniques can be applied.

- A larger number of independent observation instances is needed for nonmonotonic than for monotonic nonlinear or linear forward models for the same number of degrees of freedom in the control parameter space.
- Nonlinear monotonic forward models give rise to skewed unimodal PDFs. This property implies that data assimilation algorithms that make use of the maximum likelihood solution in theory fully account for this kind of model nonlinearity.
- The presence of model error greatly increases the possibility of capturing multiple modes in the posterior PDF when nonmonotonic nonlinearity in the model presents the potential for multiple modes to exist.
- Using a Gaussian approximation for the prior update in the presence of moderate model error has the same effect as use of larger model error and an update with the exact prior; in the case of multiple possible modes, the maximum likelihood solution in the parameter space is shifted toward the mode that is separated from the true solution. This result translates into the potential for producing a biased mean central estimate in the data assimilation solution with the Gaussian prior update, even with unbiased observations and model. This possibility does not exist only for linear models.

The results presented in this paper suggest a change in the commonly used terminology in the data assimilation literature for earth system science applications, which tends to describe forward models as quasilinear, weakly nonlinear, or strongly nonlinear. We contend, based on the results presented here, that these terms should be replaced with descriptors that characterize the model as nonlinear monotonic or nonlinear nonmonotonic, including the order of nonmonotonicity (e.g., the number of modes in the solution), if known. For example, the exponential function is highly nonlinear but monotonic, which would imply the possibility of a well-posed maximum likelihood problem, as was demonstrated for the damped oscillations model in section 4a. In contrast, the Lorenz dry cellular convection model is both nonlinear and nonmonotonic, giving rise to the possibility of multiple modes in the solution. As was demonstrated above, whether or not the inverse problem was well posed, in the sense of its ability to identify the correct parameter value from the posterior

PDF, depended on the nature of the model uncertainty, the distribution of observations in time, and the construction of the inverse problem itself.

In closing, it should be noted that this study did not examine the role of the observation operator in the inversion, nor did it consider the effect of spatial or temporal correlations between observations. Because some of the most challenging problems in inverse modeling deal with complex and often nonlinear observation operators (e.g., assimilation of clouds and precipitation) and spatially dense but nonuniform observations (e.g., high spatial and/or spectral resolution observations from polar-orbiting satellite sensors), these are topics that certainly deserve further attention. Examination of these and other issues is left for a future study.

Acknowledgments. During the course of this study, T. Vukicevic was supported by NSF Grant ATM-0514399.

REFERENCES

- Aksoy, A., A. Zhang, and J. W. Nielsen-Gammon, 2006: Ensemble-based simultaneous state and parameter estimation with MM5. *Geophys. Res. Lett.*, **33**, L12801, doi:10.1029/2006GL026186.
- Annan, J. D., D. J. Lunt, J. C. Hargreaves, and P. J. Valdes, 2005: Parameter estimation in an atmospheric GCM using the ensemble Kalman filter. *Nonlin. Processes Geophys.*, **12**, 363–371.
- Bennett, A., 1992: *Inverse Methods in Physical Oceanography*. Cambridge University Press, 346 pp.
- Braswell, B. H., B. Sacks, E. Linden, and D. S. Schimel, 2005: Estimating ecosystem process parameters by assimilation of eddy flux observations of NEE. *Global Change Biol.*, **11**, 335–355.
- Chin, T. M., M. J. Turmon, J. B. Jewell, and M. Ghil, 2007: An ensemble-based smoother with retrospectively updated weights for highly nonlinear systems. *Mon. Wea. Rev.*, **135**, 186–202.
- Evensen, G., 2006: *Data Assimilation: The Ensemble Kalman Filter*. Springer, 279 pp.
- Jazwinski, H., 1970: *Stochastic Processes and Filtering Theory*. Mathematics in Science and Engineering Series, Vol. 64, Academic Press, 376 pp.
- Kalman, R. E., 1960: A new approach to linear filtering and prediction theory. *Trans. ASME: J. Basic Eng.*, **82**, 35–45.
- Kalnay, E., 2003: *Atmospheric Modeling, Data Assimilation, and Predictability*. Cambridge University Press, 341 pp.
- Kim, S., G. L. Eyink, J. M. Restrepo, F. J. Alexander, and G. Johnson, 2003: Ensemble filtering for nonlinear dynamics. *Mon. Wea. Rev.*, **131**, 2586–2594.
- Lewis, J., S. Lakshmivarahan, and S. Dhall, 2006: *Dynamic Data Assimilation: A Least Squares Problem*. Cambridge University Press, 680 pp.
- Liu, J. S., 2001: *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, 343 pp.
- Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141.
- Mosegaard, K., and A. Tarantola, 2002: Probabilistic approach to inverse problems. *International Handbook of Earthquake and Engineering Seismology*, Vol. 1, Academic Press, 237–265.
- O'Neill, A., P.-P. Mathieu, and C. Zehner, 2004: Making the most of earth observation with data assimilation. *ESA Bull.*, **118**, 32–38.
- Rodgers, C. D., 2000: *Inverse Methods for Atmospheric Sounding: Theory and Practice*. Series on Atmospheric Oceanic and Planetary Physics, Vol. 2, World Scientific, 200 pp.
- Tarantola, A., 2005: *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 342 pp.
- Van Leeuwen, P. J., 2003: A variance-minimizing filter for large-scale applications. *Mon. Wea. Rev.*, **131**, 2071–2084.
- Vukicevic, T., B. H. Braswell, and D. Schimel, 2001: A diagnostic study of temperature controls on global terrestrial carbon exchange. *Tellus*, **53B**, 150–170.