



Data assimilation for stratified convection

Andreas Svedin,¹* Milena C. Cuéllar² and Axel Brandenburg^{3,4}

¹*Astronomy Department, Columbia University, New York, NY 10027, USA*

²*CUNY - LaGuardia Community College, 31-10 Thomson Avenue, Long Island City, NY 11101, USA*

³*Nordita, KTH Royal Institute of Technology and Stockholm University, Roslagstullsbacken 23, SE-10691 Stockholm, Sweden*

⁴*Department of Astronomy, AlbaNova University Center, Stockholm University, SE-10691 Stockholm, Sweden*

Accepted 2013 May 20. Received 2013 May 19; in original form 2012 July 31

ABSTRACT

We show how the 3DVAR data assimilation methodology can be used in the astrophysical context of a two-dimensional convection flow. We study the way in which this variational approach finds best estimates of the current state of the flow from a weighted average of model states and observations. We use numerical simulations to generate synthetic observations of a vertical two-dimensional slice of the outer part of the solar convection zone for varying noise levels, and implement 3DVAR when the covariance matrices are diagonal and proportional to the identity matrix. Our simulation results demonstrate the capability of 3DVAR to produce error estimates of system states that can be more than two orders of magnitude below the original noise level present in the observations. This work illustrates the importance of applying data to obtain accurate model estimates given a set of observations. It also exemplifies how data assimilation techniques can be applied to simulations of stratified convection.

Key words: chaos – convection – MHD – methods: numerical – methods: statistical – Sun: general.

1 INTRODUCTION

When using models to describe the temporal evolution of observed complex systems we are confronted with a number of challenges. An immediate difficulty in dealing with this issue is that we generally do not know in full detail the current state of the system or the initial conditions that are to be used. The lack of such information prevents us from keeping a model-based simulation in step with the behaviour of the observed system.

Data assimilation techniques offer a means to address such challenges for complex systems by keeping a computer simulation (i.e. model) in synchronization with observations of the system it represents. They provide a general framework for simultaneously comparing, combining and evaluating observations of physical systems and output from computer simulations.

The methods used in data assimilation have been developed over several decades, primarily in meteorology and oceanography for the prediction of future behaviour. Data assimilation is used daily in operational weather prediction (Bengtsson, Ghil & Källén 1981; Kalnay 2003), in climate forecasts (Palmer & Hagedorn 2006), and it was even used to correct the path of the Apollo spacecraft during the first moon landings (Cipra 1993). There is a large and growing body of literature, including several monographs (e.g. Daley 1993; Kalnay 2003; Wunsch 2006) and papers discussing its theoretical foundations (e.g. Lorenc 1981; Lorenc 1986; Le Dimet & Talagrand 1986; Ghil 1989). Astrophysical data assimilation has recently been

discussed by Brun (2007), in the context of space weather and solar cycle prediction (Dikpati 2007; Choudhuri, Chatterjee & Jiang 2007; Kitiashvili & Kosovichev 2008), as well as in dynamo models (Jouve, Brun & Talagrand 2011).

Here we focus on the three-dimensional variational (3DVAR) data assimilation technique, which is a sequential approach (Daley 1993). It produces updates of the current state of a model simulation at times when system observations are available. Propagation of model states between updates is performed using a model run initiated at the latest state estimate. An extension of 3DVAR to implicitly incorporate dynamical information is known as four-dimensional variational (4DVAR) data assimilation. 3DVAR produces estimates only at times when an observation and a model estimate are available. At other times, model trajectories fill the gaps between the observations. The 4DVAR estimate is over an interval in time during which all observations in the interval have been accounted for, and the model dynamics is used to propagate information within the interval.

State estimates produced by 3DVAR are optimal provided that the model is linear and the error statistics are Gaussian, stationary and correctly specified. In other words, 3DVAR states are best linear unbiased estimates, where *best* and *optimal* refer to the lowest possible mean squared error of the estimate (Kalman 1960; Talagrand 1997).

For non-linear models, error statistics may become non-Gaussian even when the initial distribution is normal and 3DVAR (or 4DVAR) estimates are no longer unbiased. In this case, data assimilation techniques are challenged by the fact that actual applications are

*E-mail: aos2112@columbia.edu

typically based on non-linear processes (Pires, Vautard & Talagrand 1996). Specifically, the states exhibited by real systems under observation will diverge from those predicted by a model simulation. This has two principal causes (Palmer & Hagedorn 2006): observational error and sensitivity to initial conditions. The first of these is a result of what may be called noise. Because its statistical character may not be known, we may need to make some assumptions about its properties. The second source of error occurs in many complex systems and is referred to as chaotic behaviour. This has been known for some time, but only in recent decades has serious progress in its understanding been possible. The sensitivity of the model to initial conditions limits how far into the future predictions can be made (Lorenz 1993). Two examples of other data assimilation challenges are the mismatch between spatial locations of observations and grid positions of the model, and model variables not corresponding to observed quantities.

Despite the challenges and some open questions, 3DVAR is widely used in the oceanographic and meteorological communities, and would make a good candidate method to explore in the context of astrophysical flows.

2 STRATIFIED CONVECTION MODEL

We are motivated to use data assimilation techniques in the context of stratified convection as a path to obtain predictions of solar subsurface weather events, namely of the flow structure beneath the surface. The ability to anticipate the possibility of violent events on the solar surface, such as coronal mass ejections that affect the space weather and the dynamics of the Earth's magnetosphere, is important (see Ionidis, Zhao & Kosovichev 2011). The idea is to use a model of solar subsurface convection, ultimately involving the magnetic fields that give rise to surface activity such as sunspots, solar flares, and coronal mass ejections, although current advances in that direction are still at a preliminary stage (Warnecke, Brandenburg & Mitra 2011). However, once such models are able to reproduce sufficient details of solar activity, it will be important to synchronize a given model with daily observations to be able to use it for predictions.

As a proof of concept, we design a data assimilation experiment to test the implementation of 3DVAR for the `PENCIL CODE`, a public-domain code of high order (sixth order in space and third order in time) for solving the hydrodynamic equations (Brandenburg & Dobler 2002).¹ We consider here a simple two-dimensional convection model representing the turbulent flows of stars with outer convection zones. In our experiment, *synthetic observations* are generated by adding noise to the output from our model. These observations are then processed by 3DVAR to produce an *analysis*. An analysis is an estimation of the *unknown* state of a system in terms of model variables (Lorenz 1986; Talagrand 1997).

Our implementation is general and in the future could be used for other problems that can be addressed with the `PENCIL CODE`. In this work we assume the model to be *ideal* and to reproduce faithfully the features present in the observations. In real-world applications, the models are far from ideal, and imperfections and uncertainties related to the model are always present. Ideally, we would like to be able to account for some portion of those unknowns by using data assimilation techniques.

We use the sample `2d/conv-slab-MLT` of the `PENCIL CODE` (revision `r14696` and later). This sample simulates a vertical

two-dimensional slice of the outer part of a stellar convection zone. In particular, we use it to simulate convection at low resolution, 64×64 , at a Rayleigh number of 8×10^5 (Dintrans et al. 2005), and at a Reynolds number of ~ 30 . The basic setup is similar to that described in Brandenburg et al. (2005) and for other earlier models (Hurlburt, Toomre & Massaguer 1986; Brandenburg et al. 1996), consisting of a convectively unstable layer sandwiched between two stable layers.

The simulated vertical two-dimensional slice of the outer part of the solar convection zone has a mean field velocity of $u_{\text{rms}} = 0.08$, the wavenumber of the energy-carrying eddies is $k_f = 2\pi/d$ for a depth of the unstable layer $d = 1$, and therefore the correlation time $\tau_{\text{cor}} = (u_{\text{rms}}k_f)^{-1}$ is ~ 2 . Starting from an initial velocity field of perturbations with amplitude of 3×10^{-4} times the average sound speed, convective motion is generated without having to introduce any other forcing. This model is chosen to illustrate 3DVAR in an astrophysical context for sufficiently complex behaviour without including any stochastic elements.

3 DATA ASSIMILATION SETUP

The 3DVAR scheme was developed in the meteorological community to improve model-based weather prediction in the face of observational and modelling uncertainties. It was formulated in a unified Bayesian framework by Lorenz (1986). 3DVAR produces an update of the current state of the system at times when observations are available, which in turn can be used as a new initial condition to propagate the model forwards to the time when the next observation is available.

We can use 3DVAR as a black box along with a low-resolution simulation to assimilate many data points at low computational cost on a laptop computer. For example, a typical model run for a 64×64 two-dimensional convection field over a time interval of 300 time units (corresponding to about 150 correlation times or about 1 d if applied to the top of the solar convection zone where typical velocities are 1 km s^{-1} and typical length scales are 500 km) takes about 15 min on a laptop computer. A single data assimilation correction takes 4 min. We expect that an optimized implementation could be much more efficient. 3DVAR minimizes the sum of the squared differences between the current model state and the observations to find a solution that is a compromise between these two estimates of the true state.

It is important to realize that, in real problems, the true state is available only through noisy observations of the system. We have the ideal case in which there is no model uncertainty and the only source of uncertainty is in the observations. In this way, we can assess how far/close the model state is to the true state of the system. The key is to generate a known true state against which the estimated state obtained via data assimilation can be verified.

In our *twin-experiment* (Bengtsson et al. 1981) we select two different initial fields to run the `PENCIL CODE` simulation. One of these initial fields represents the unknown true initial state of the system. The other initial field represents what might be, in practice, a good approximation, or guess, of the initial state of the system.

The initial field chosen to represent the true initial state of the system initializes a model run corresponding to the original state of the system. This is to be used as a reference trajectory or *control*. The other initial field is used to initialize two different runs: one *free* model run and another that will become the assimilated trajectory or *analysis*. The analysis is a collection of segments of model trajectories initialized at the 3DVAR corrections made at all times when the observations are available. The model state before

¹ <http://pencil-code.googlecode.com/>

correction is called the *background state*, and the updated estimate is called the *analysis*.

Comparing the free and the control run gives us a measure of the sensitivity to initial conditions of our model; that is, it shows how similar initial conditions diverge in time. Similarly, a comparison of the free run and the analysis represents the effect of the data assimilation procedure over a trajectory, where both start from the same initial field. If the assimilation of the second set of initial conditions is effective it will bring the analysis ‘closer’ to the control run.

We generate synthetic observations by adding independent and identically distributed noise to the horizontal velocity field at all grid points of the control run. These synthetic data are considered to be our experimental observations, which in turn will be used to update the analysis. As explained in the next section, 3DVAR requires both a model state and an observation to update the analysis at each assimilation.

4 3DVAR AND THE WEIGHT FACTOR

The 3DVAR technique finds a model state \mathbf{x} that agrees with the current state of the observed system given the information available in the observations and the model. Specifically, we minimize the weighted average of the residues for both observations \mathbf{y}_0 and model states \mathbf{x}_b at a specific time t to find an optimal solution. This is expressed in the 3DVAR cost function (Lorenz 1986)

$$J(\mathbf{x}) = \frac{1}{2} [\mathbf{x} - \mathbf{x}_b]^\top \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_b] + \frac{1}{2} [\mathbf{y}_0 - \mathbf{H}(\mathbf{x})]^\top \mathbf{R}^{-1} [\mathbf{y}_0 - \mathbf{H}(\mathbf{x})], \quad (1)$$

where \mathbf{x}_b is the model state – traditionally called the *background state* – and \mathbf{y}_0 is the observed state. The background covariance matrix is defined as

$$\mathbf{B} = \left\langle (\delta \mathbf{u}_x^\Delta - \langle \delta \mathbf{u}_x^\Delta \rangle) (\delta \mathbf{u}_x^\Delta - \langle \delta \mathbf{u}_x^\Delta \rangle)^\top \right\rangle, \quad (2)$$

with averages denoted by $\langle \cdot \rangle$, and where $\delta \mathbf{u}_x^\Delta$ is defined as the difference between the background state \mathbf{x}_b and the reference state \mathbf{x}_r :

$$\delta \mathbf{u}_x^\Delta = \mathbf{x}_b - \mathbf{x}_r. \quad (3)$$

The observational covariance matrix is in a similar way defined as

$$\mathbf{R} = \left\langle (\delta \mathbf{u}_x^0 - \langle \delta \mathbf{u}_x^0 \rangle) (\delta \mathbf{u}_x^0 - \langle \delta \mathbf{u}_x^0 \rangle)^\top \right\rangle, \quad (4)$$

where $\delta \mathbf{u}_x^0$ is defined as the difference between the observations \mathbf{y}_0 and the corresponding reference state $\mathbf{H}(\mathbf{x}_r)$:

$$\delta \mathbf{u}_x^0 = \mathbf{y}_0 - \mathbf{H}(\mathbf{x}_r), \quad (5)$$

where the observation operator \mathbf{H} has been used to project the model variables onto observables.

As a result of minimizing $J(\mathbf{x})$ in equation (1) we obtain the analysis \mathbf{x}_a corresponding to the best estimate of the current state of the system. After the estimate is generated, the model is integrated forwards using the analysis as the initial condition to the next time an observation is available. Synthetic observations are denoted by \mathbf{y}_0 , and contain normally distributed noise of amplitude σ_R , proportional to the maximum amplitude of the full two-dimensional vertical velocity field.

For example, a noise level of 1 per cent corresponds to $\sigma_R = 0.01$ times the peak-to-peak amplitude for a normalized field or to $\sigma_R = 6 \times 10^{-3}$ for an unnormalized field. A realistic noise level

for observations in the vertical component of the velocity at the top 3.5 to 5.5 Mm of the convection zone is one between 0.007 and 1 per cent of the peak-to-peak amplitude (Švanda et al. 2011).

The selection and construction of the observational and background covariance matrices (\mathbf{R} and \mathbf{B}) are of great interest in data assimilation (Bannister 2008a,b). In our case, the observational noise is not correlated in space, and we neglect spatial correlations between model states, making the off-diagonal components of the \mathbf{B} and \mathbf{R} matrices vanish. By doing this, we can set these matrices to be proportional to identity matrices (referred to as scalar variances), $R_{ij} = \delta_{ij} \sigma_R^2$ and $B_{ij} = \delta_{ij} \sigma_B^2$. Without these spatial correlations 3DVAR generates an analysis that is, in general, less smooth over the two-dimensional domain. In more sophisticated formulations of equation (1), the form of \mathbf{B} can also include physical constraints to processes not resolved in the model (Dobricic & Pinardi 2008).

In turn, we set the observation operator to be $H_{ij} = \delta_{ij}$. This means that we assume that the observations cover all grid points in the model domain; that is, that the observables and model variables belong to the same space. In other words, the system and model are the same. In more realistic implementations of 3DVAR, \mathbf{H} is typically a computer algorithm that cannot be expressed explicitly as a matrix owing to its non-linear nature (Dobricic & Pinardi 2008).

After these assumptions, we can translate the cost function (equation 1) to

$$J(\mathbf{x}) = w (\mathbf{x} - \mathbf{x}_b)^2 + (\mathbf{x} - \mathbf{y}_0)^2, \quad (6)$$

where w is the ratio of the scalar variances corresponding to the observed and background states

$$w = (\sigma_R / \sigma_B)^2. \quad (7)$$

In this setting, to find the state vector \mathbf{x}_a that minimizes (6), we use Powell minimization (Press et al. 1992). The advantage of this procedure is that no information about the gradient or adjoint of the model is needed. This would be required for more efficient optimization schemes such as the quasi-Newton or the conjugate gradient method (Press et al. 1992).

The coefficient w in equation (6) behaves as a *weight factor* in the optimization process and will be referred to as such. Solving $\nabla J(\mathbf{x}) = \mathbf{0}$ yields

$$w(\mathbf{x} - \mathbf{x}_b) + (\mathbf{x} - \mathbf{y}_0) = \mathbf{0}, \quad (8)$$

and the optimal state of the model, the analysis $\mathbf{x}_a \equiv \mathbf{x}$, that represents the systems given the current observations is found by solving equation (8) for \mathbf{x} .

$$\mathbf{x}_a = \frac{1}{w+1} \mathbf{y}_0 + \frac{w}{w+1} \mathbf{x}_b \equiv \mathbf{x}_b + \frac{\mathbf{y}_0 - \mathbf{x}_b}{w+1}. \quad (9)$$

Equation (9) reveals how the contribution of the background states \mathbf{x}_b and observations \mathbf{y}_0 affects the analysis \mathbf{x}_a in terms of the weight factor w . Fig. 1 clearly shows the result of this process. At times when observations (grey open diamonds) are available, 3DVAR performs a correction, given by equation (6), to the analysis (black plus signs), and from which a segment of background states (dotted curve) is initialized and run up to the next assimilation time. The analysis is the background states for times different from the assimilation times, and the corrected values obtained at assimilation times.

Given a model and a fixed set of observations, equation (9) help us to understand the effects of the weight factor in the resulting analysis \mathbf{x}_a . The case $w < 1$, or $\sigma_R < \sigma_B$, corresponds to the observational uncertainty being smaller than the model uncertainty. In

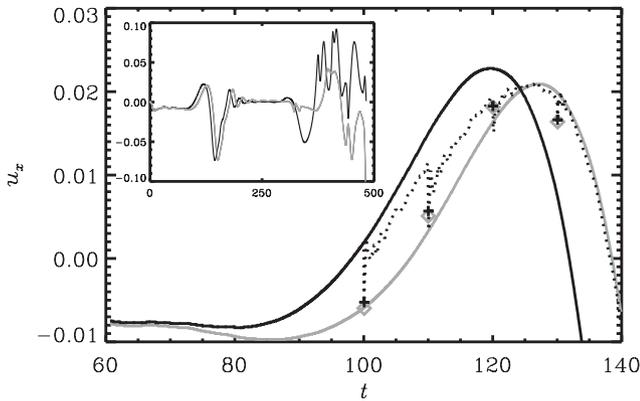


Figure 1. Temporal evolution of the horizontal velocity at a certain midpoint of a two-dimensional convection field for $t \in [0, 500]$. The control, free run and analysis correspond to the grey, black and dotted curves. Grey diamonds represent observations, and plus signs ‘+’ represent corrections (\mathbf{x}_a) made by 3DVAR, both at assimilation time. The corrections ‘+’ take the estimate (dotted line) towards the true state of the observed system (grey line).

other words, weight is given to the observations because small w allows the distance $|\mathbf{x}_a - \mathbf{x}_b|$ to grow without making large contributions to the cost function. In contrast, having $w > 1$, or $\sigma_R > \sigma_B$, favours model states corresponding to more uncertainty related to the observations than to the model.

The next section presents and describes the results of our numerical experiments. We study the ‘quality’ of the 3DVAR analyses by varying the value of the weight factor w , when applying the procedure to two sets of observations with different noise levels.

We calculate the correlation time to be approximately two using our simulation setup parameters (see Section 2).

We choose an assimilation time large enough to let the oscillations propagate over the two-dimensional field but still small enough to be able to capture the smaller-scale dynamics. In each 3DVAR experiment, data assimilation corrections are made every 10 time units (which corresponds to about 5 correlation times). We found no fundamental difference when using assimilation times between 5 and 20 time units.

5 RESULTS

We generate analyses using 3DVAR for each of the weight factors $w = \{0.1, 1, 10\}$. For a fixed value of w we generate an analysis that is the result of assimilating one set of observations with 1, 2, 5 or 10 per cent noise levels for two different fields of initial conditions. The resulting horizontal velocity, u_x , at the midpoint of the upper right quadrant of the two-dimensional domain, is plotted in Fig. 2 for $w = 0$ in the upper panel and for $w = 0.5$ in the lower panel, and in Fig. 3 for $w = 1$ in the upper panel and for $w = 10$ in the lower panel.

Note that the grey and dashed lines are the same in all panels and they represent the reference states of the system (control) and the corresponding free run of the model. Observations are plotted with grey diamond symbols. Black plus marks (+) are used for the corrections calculated at the time of assimilation by minimizing expression (6). Between assimilations, the analysis is the segment of model trajectory (dotted segments) initialized at the corrected state ‘+’, as seen in detail in all insets in Figs 2 and 3.

From equation (9) and these insets, the amplitude of the correction made in each case is clearly illustrated for each value of the weight

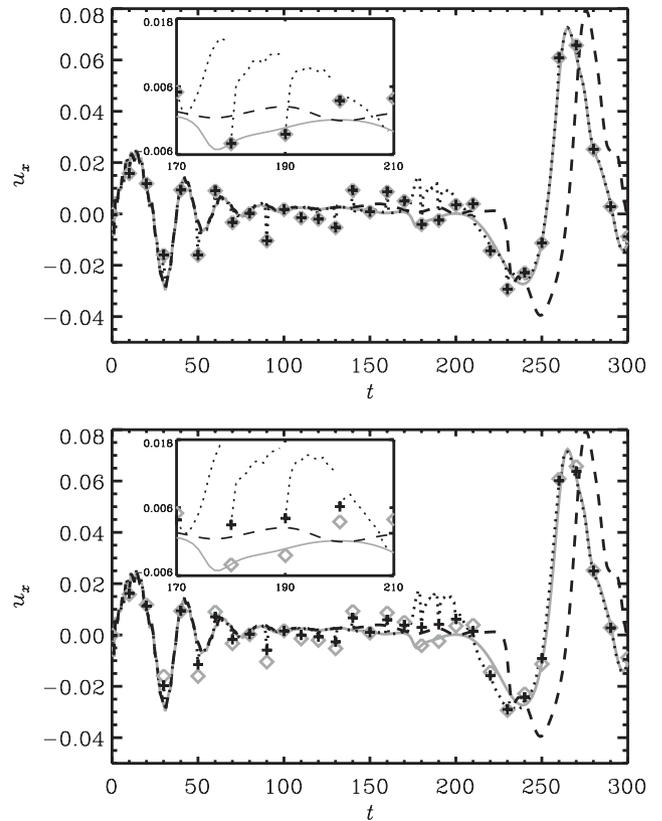


Figure 2. Data assimilation run over observations marked with grey diamonds with 1 per cent noise for $w = 0$ (upper panel) and $w = 0.5$ (lower panel). Black plus signs mark the 3DVAR corrections at assimilation times. The grey, dotted, and dashed curves correspond to the control, analysis, and free trajectories, respectively. Both panels include zoom-ins for $170 \leq t \leq 210$.

factor w . This amplitude is measured by the gap between the ‘+’ and the last dot of the previous dotted segment (background states).

For the trivial case of setting $w = 0$, the second term in (9) is zero, and the best estimate of the current state of the system is given by $\mathbf{x}_a = \mathbf{y}_0$, as seen in the upper panel of Fig. 2. The correction (‘+’) is ‘pulled’ from the background state (dotted curve) to the observation (‘◇’) at the time of assimilation. This specifically corresponds to replacing the background state with the observed values at the observed points (i.e. at all grid points of the model).

For any other value of $w > 0$ the correction is in between the observation (‘◇’) and the last state of the previous background segment (dotted lines); see insets in Figs 2 and 3.

For $0 < w < 1$, the optimal value of the cost function (6) is a factor $1/w$ closer to the observations, \mathbf{y}_o , than to the background state \mathbf{x}_b ; see equation (9). The lower panel of Fig. 2 shows the results for $w = 0.5$, where the corrections (‘+’) fall closer to the observations than to the end of the last background states. This figure shows that the analysis follows the control trajectory closely (solid grey line). Note that even if the corrections are large, for example at $t = 50$ or $t = 90$, the analysis quickly relaxes to the control state.

In the case $w = 1$, equal weights are given to model states and observations. The optimal value of \mathbf{x}_a is the average of \mathbf{y}_0 and \mathbf{x}_b , from (9). No preference is given to any estimate, and the midpoint is the optimal choice for \mathbf{x}_a as seen in the upper panel of Fig. 3.

When $w > 1$, the optimization of equation (6) will favour model states rather than observations as follows from equation (9). The analysis at assimilation times is closer to the model states than to the

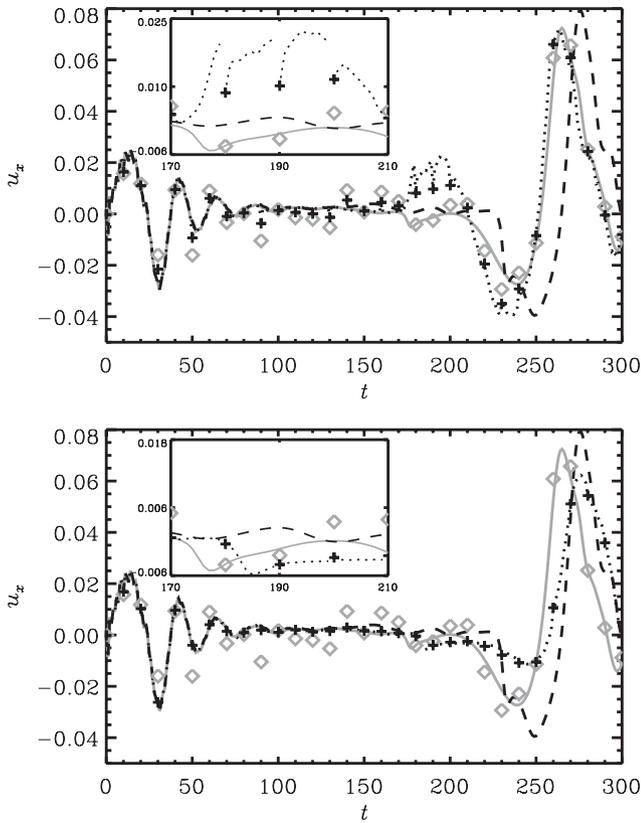


Figure 3. Data assimilation run over observations marked with grey diamonds with 1 per cent noise for $w = 1$ (upper panel) and $w = 10$ (lower panel). Black plus signs mark the 3DVAR corrections at assimilation times. The grey, dotted, and dashed curves correspond to the control, analysis, and free trajectories, respectively. Both panels include zoom-ins for $170 \leq t \leq 210$.

observations by a factor w . In the lower panel of Fig. 3, the resulting trajectories are plotted specifically for $w = 10$. We observe from this plot and from results at other locations of the two-dimensional domain that, for $w > 1$, the estimates of the original state of the system are biased towards the background states.

In this simplified experiment, we can see the great importance of the weight factor w in 3DVAR: this emphasizes out how crucial the construction of the covariance matrices \mathbf{B} and \mathbf{R} is for optimal results. One of the motivations for our choice of \mathbf{R} and \mathbf{B} to be proportional to identity matrices is to set a baseline from which we can illustrate in a simplified analytical way the inner workings of 3DVAR. It can be hard to see how the different components interact to create a result when more sophisticated choices of \mathbf{R} and \mathbf{B} are used.

We note that, in Figs 2 and 3, we chose the interval $t \in [170, 210]$ as an example of an interval in which 3DVAR does not perform very well. In that range, 3DVAR systematically pulls the analysis away from the control – considered here as the original system trajectory. Ascertaining the reason for this behaviour requires further study, but it is worth noticing that the high performance of the data assimilation returns around $t = 200$. The case $w \geq 10$, in the lower panel of Fig. 3, actually performs better in this interval.

Consistently, we observe that the analysis is on average closer to the control trajectory than to the observations for all values $w \leq 1$. As noted, exceptions are observed for larger values of the weight factor and during the interval shown in the insets of Figs 2 and 3.

Table 1. Measures of variability (in 10^{-6}) for the distance between the control and observations with noise levels of 1, 2, 5 and 10 per cent and the free trajectory. See text for a description of the notation.

Noise level	$(\delta \mathbf{u}_x^0)_{\text{mid}}$	$\langle\langle (\delta \mathbf{u}_x^0)^2 \rangle\rangle_1$	$\langle\langle (\delta \mathbf{u}_x^0)^2 \rangle\rangle_2$
Free	1100	530	3000
1 per cent	27	36	36
2 per cent	108	145	143
5 per cent	888	899	900
10 per cent	3553	3596	3602

Tables 1 and 2 present several measures of variation of the output from the simulations of our twin experiment. For each simulation, we calculate the variance of the distances between the control and the observations (Table 1) and between the control and the analysis (Table 2) over the data assimilation window.

Specifically, Table 1 shows the variance of the distance between the control trajectory and the free trajectory (first row), and the noisy observations for 1 per cent (second row) and 2 per cent (third row) noise levels, followed by 5 per cent and finally 10 per cent. The second column of Table 1 presents the variance of the values at the midpoint of the field denoted by $(\delta \mathbf{u}_x^0)_{\text{mid}}$; these values measure the variability of the local behaviour. The third and fourth columns of this table shows the averaged variance over the whole vertical two-dimensional field, for the first and the second half of the assimilation window respectively. This variance is denoted by the inner angle brackets, $\langle\langle (\delta^O \mathbf{u}_x)^2 \rangle\rangle_T$, where $T = 1, 2$ correspond to the assimilation window intervals $t \in [1, 150]$ and $t \in [151, 300]$, respectively.

Table 2. Measures of variability (in 10^{-6}) for the distance between the control and the analysis for several values of the w of the x -component of the velocity, \mathbf{u}_x . The italic values are larger than the corresponding noise level. See text for a description of the notation.

w	$(\delta \mathbf{u}_x^A)_{\text{mid}}$	$\langle\langle (\delta \mathbf{u}_x^A)^2 \rangle\rangle_1$	$\langle\langle (\delta \mathbf{u}_x^A)^2 \rangle\rangle_2$
Noise level 1 per cent			
0	0.43	0.34	3.83
0.1	0.47	0.35	4.40
1	2.96	4.09	<i>130.92</i>
10	<i>271.65</i>	<i>506.57</i>	<i>1534.65</i>
Noise level 2 per cent			
0	1.84	1.29	15.45
0.1	1.99	1.34	17.86
1	25.30	12.24	<i>751.81</i>
10	<i>888.62</i>	<i>718.40</i>	<i>2823.14</i>
Noise level 5 per cent			
0	5.17	6.57	15.47
0.1	4.54	5.96	17.54
1	8.10	6.37	60.48
10	607.71	<i>818.45</i>	<i>3003.26</i>
Noise level 10 per cent			
0	22.36	27.45	70.99
0.1	20.13	24.63	79.58
1	40.21	26.19	237.57
10	1087.08	2676.32	<i>4775.14</i>

Note that from the values in Table 1, the free run is one or two orders of magnitude further away from the original state of the system (control) than the 1 and 2 per cent noisy observations. The large difference between the second and third columns reflects how the free run is diverging from the control run over two different time intervals. Table 1 is the baseline from which we measure the performance of 3DVAR when estimating the state of the system from noisy observations or the free trajectory (i.e. trajectories corresponding to the 0 per cent noise level).

When assessing the performance of 3DVAR for different values of w and noise levels, we look for variability measures between the control and the analysis lower than the levels set by the free run and the noisy observations. As can be seen in Table 1, all variance measures for the free trajectory are larger than the corresponding variance values for all w and noise levels up to 5 per cent. This means that performing 3DVAR data assimilation is more effective at estimating the original state of the system than just using as an estimate the trajectory initialized with a very close initial field.

The variance of the distance between the two initial velocity fields is $\sim 0.18 \times 10^{-6}$. Table 2 shows the variance of the distance between the control run and the analysis at the midpoint of the field in the second column, denoted by $\langle (\delta u_x^A)_{\text{mid}}^2 \rangle$. In the third and fourth columns of the table, we show the averaged variance over the whole vertical two-dimensional field for the first and the second half of the assimilation window, denoted by $\langle \langle (\delta^A \mathbf{u}_x)^2 \rangle \rangle_T$, for $T = 1, 2$, respectively.

From the values presented in both tables, we observe that for noise levels below 10 per cent and $w < 10$, the variance of the distance between the control and the analysis is smaller – on average – than the corresponding variance of the distance between the control and the observations at the same noise level (see Table 1). This means that over the two-dimensional domain, on both local and global scales, 3DVAR is effectively reducing noise and accounting for sensitivity to initial conditions; that is, it is estimating a value for the horizontal velocity closer to the original state of the system (control) than to the trajectory generated using a guess of the initial state of the system (free run).

In contrast, for $w = 10$ and all noise levels at both global and local scales, the free run and the analysis are on average consistently further away from the control than the original noisy observations (corresponding values in Table 1). This is also observed, for example, at the specific location of the two-dimensional domain shown in Figs 2 and 3.

More sophisticated choices for \mathbf{R} and \mathbf{B} , which inform the cost function more realistically (6) about spatial correlations over the field, might generate an improved and more consistent performance.

Furthermore, in Fig. 4 we plot on a semi-logarithmic scale the variance of the point-by-point distances of the two-dimensional vertical slices between the control and analysis (grey solid circles), the control and the observations (grey diamonds), and the control and the free run (black solid circles) for all $t \in [1, 300]$. The black plus signs mark the variance of distances between the control and analysis fields at assimilation times. From the top panel to the bottom panel, plots correspond to $w = 0.1, 1, 10$, for a 1 per cent noise level. Black dotted and grey diamond curves are the same for all panels in the figure.

It is important to note that the values plotted in Fig. 4 are not the running variance of the distance between control trajectories and the other relevant trajectories but the variance of the distances between the two-dimensional fields at each step.

We observe in Fig. 4 that for all values of w and for $t \in [0, 10]$ the analysis and free runs are good estimators of the control

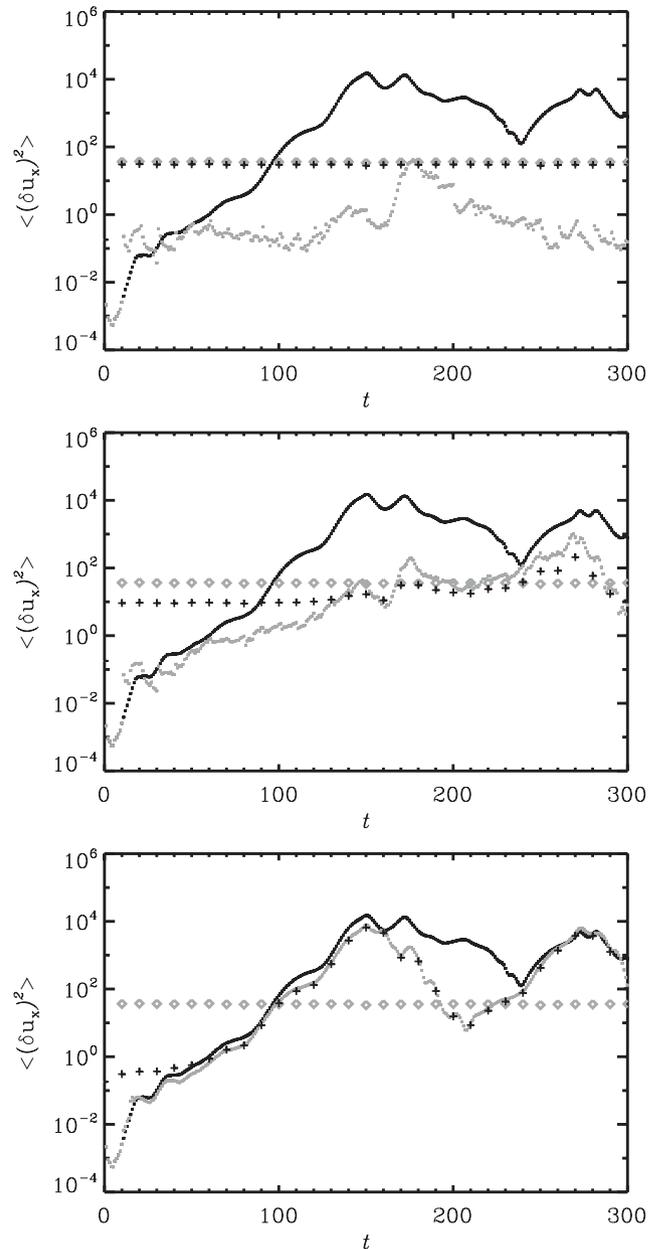


Figure 4. Semi-logarithmic plot of the variance of the point-by-point distances of the two-dimensional vertical slices between the control and the analysis (grey solid circles, $\langle (\delta^A \mathbf{u}_x)^2 \rangle$), between the control and the observations (grey diamonds, $\langle (\delta^O \mathbf{u}_x)^2 \rangle$), and between the control and the free run (black solid circles, $\langle (\delta^F \mathbf{u}_x)^2 \rangle$) for all $t \in [1, 300]$. The black plus signs correspond to $\langle (\delta^A \mathbf{u}_x)^2 \rangle$ at assimilation times, after the correction is made. All quantities are scaled by 10^6 . From the top to bottom panels, $w = 0.1, 1, 10$.

trajectory. In addition, for all values of w and $t \in [20, 100]$, the free run and the analysis variance with respect to the control run are below the variance of the noise level (the grey curve is below the black curve), showing that for a 1 per cent noise level 3DVAR produces a better estimate than the free run. Only for $w = 0.1$ is the 3DVAR analysis a better estimate of the state of the system than the free run or control run for all values in $t \in [30, 300]$. Note that there is an increase in the error of the estimate from $t \in [160, 190]$ with respect to the control, which corresponds to the interval shown in the insets in Figs 2 and 3. Otherwise, the error of the estimate

exceeds the noise level for $w = 1$ at $t \in [170, 300]$, and for $w = 10$ at $t \in [190, 230]$.

The noise reduction of 3DVAR is clearly seen in the top panel of Fig. 4. At time $t = 300$ the grey estimate has a value of 0.15×10^{-6} , a factor of ~ 230 below the noise level of 36×10^{-6} . This should be compared with the variance of the distance between the two initial velocity fields of 0.18×10^{-6} . The error of the estimate is also below the noise level for most cases ($w < 10$ per cent) presented in Table 2. The important point is that our estimates not only are ‘nudged’ towards the observations, but are steered towards the true unknown state of the system, effectively reducing the noise of the observations. The magnitude of this noise reduction, found in this preliminary study, should not be expected in real-world applications. The quality of the estimates ultimately depends on the quality of the model and the observations.

Fig. 4 also shows how far apart – on average – the control and the free run become as time increases. The increase in the amplitude of the variance (the black curve) from $t = 0$ to $t = 300$ is observed to be up to six or seven orders of magnitude.

We again see evidence of how the 3DVAR correction (black plus signs) lands closer to the observations (diamonds) for lower values of w , and how the background states start to move towards the grey model trajectory for higher values of w , as observed in earlier figures. The effect of $w < 1$ is to pull the background states back closer to the observations, as the grey dotted curve is below the corrections (plus signs).

Recall that the weight factor is $w = (\sigma_R/\sigma_B)^2$, and it is interpreted here as a measure of the relative confidence given to either the observations or the model. High values of w reflect a higher trust in the observations than in the model representation of the system, and the opposite is the case for low values of w . In this simple setting, σ_B relates to the sensitivity to initial condition rather than to model deficiencies. As noted earlier, the estimate is sensitive to the choice of σ_B , and a more sophisticated choice can include additional components that might help to alleviate some of the model deficiencies.

In the particular case where the covariance matrices are assumed to be scalar and the weight factor is small, we can conclude that the effect of trusting the observations more than the model states ($\sigma_R < \sigma_B$) provides a closer estimate of the original state of the system than just generating a trajectory close to the initial condition. This means that 3DVAR is successful at finding an optimal state estimator in the limit of small observational noise, small model uncertainty related only to sensitivity to initial conditions, and scalar covariance matrices, \mathbf{R} and \mathbf{B} . Our results emphasize the importance of choosing more sophisticated covariance matrices (\mathbf{R} and \mathbf{B}) to better reflect the known uncertainty sources in the data assimilation problem of interest.

6 DISCUSSION

We have presented an idealized case in which the model and system are the same: a computer simulation is used both to generate synthetic observations and as the model required for the data assimilation procedure. In this way, we can assess the closeness of the model estimate to the true state of the system. The key is to have access to the true states, which we can use to verify and evaluate estimates obtained using data assimilation.

We used a simplified formulation of the 3DVAR data assimilation technique in terms of the weight factor: $w = (\sigma_R/\sigma_B)^2$. The weight factor defines the contribution of the model states, which contain propagated information from previous observations, and the current

observation to make a state estimate. The formulation of 3DVAR used here is achieved by reducing the covariance matrices, \mathbf{R} and \mathbf{B} , to diagonal matrices, and the observation operator, \mathbf{H} , to the identity. This selection corresponds to neglecting all spatial correlations between model states over the two-dimensional domain in addition to one-to-one correspondences between system observables and model variables. In this way, we clearly separate the contribution of observations and model states to the estimated state, as seen in equations (6) and (9). We solved equation (6) using a minimization scheme and found that it agrees with the analytic expression (9).

It is less obvious to see how the different components would interact to create an estimate of the original state of the system when more sophisticated choices of covariance matrices are made to fully represent uncertainties and spatial correlations. In those cases, we would have to think about the optimal combination in analogy to equation (9), in terms of a generalization of the weight factor w as a *weight matrix*, $\mathbf{W} = \mathbf{R}\mathbf{B}^{-1}$. In this analogy, model states and observations will be projected by the matrices $\mathbf{W}[\mathbf{W} + \mathbf{I}]^{-1}$ and $[\mathbf{W} + \mathbf{I}]^{-1}$, respectively, onto the analysis \mathbf{x}_a . Here, \mathbf{I} is the identity matrix.

In general, we can say that to understand the 3DVAR algorithm it is important to look at the weight factor, particularly in the limit where \mathbf{W} is assumed to be the scalar w . Consistently we observe the error between the state estimate and the original state of the system to be below the noise level when more weight is given to the observations than to the model states and the noise level is below 10 per cent. When the contribution from the model states is larger than the contribution from the observations, we note that the error eventually becomes larger than the noise level; see the case for $w = 10$.

We note in Figs 2, 3 and 4 that 3DVAR under-performs for $t \in [160, 230]$ both locally and globally for all values of w . Further study of the simulation is needed to account for this atypical behaviour. Minute differences in initial conditions generate different time evolutions for the different runs, as is expected for non-linear systems. This is illustrated by the black curves in Fig. 4 that present how the variance of the distance between the two initial conditions grows over the time interval. It can also be seen in Fig. 1, where the grey and black curves (which started with close initial condition) are very different at later times. On the other hand, a large correction made by 3DVAR, for example the black plus signs at times 50 and 90 in Fig. 2, does not appear to have a strong effect. The model run that starts at these far away estimates converges almost instantly back to the control trajectory at that time.

These features, reminiscent of chaotic behaviour, can be understood in terms of attracting sets, where small changes in the initial conditions generate a different time evolution on the attracting set. A large correction probably takes us outside the attracting set, and the solution rapidly falls back when the model is integrated forwards. These features can also be interpreted in more physical terms, as a large correction takes us to states that are not consistent with conservation laws and other physical constraints. The system would then rapidly be forced back on to a more physical state. This dual picture, using both physical and mathematical intuition, helps us to understand the observed behaviour.

The 3DVAR methodology is optimal, provided that it is correctly specified, for linear models with stationary Gaussian error statistics. Very few models in astrophysics have these properties. The validity of variational methods outside the linear or weakly non-linear case is unclear, but if the assimilation is frequent enough the behaviour might be close to linear. Higher assimilation frequency during the interval $t \in [170, 230]$ might, for example, give a better result.

One possible way to better account for non-Gaussian error statistics that change in time would be to use a hybrid-3DVAR approach (Wang, Snyder & Hamill 2007), in which the background covariances are explicitly modelled by an ensemble.

The indication of chaotic properties and attracting state space sets invites the use of other data assimilation methods that explicitly take these properties into account (Judd & Smith 2001; Judd et al. 2008). These might be more applicable to non-linear astrophysical processes. In this work we used an ideal model; that is, the system and the model are one and the same. There will always be some limitations to the modelling of a real system that might prove problematic. In general, it is more important for models used for the prediction of real systems to be reasonably realistic than just to be close to the observations.

For non-linear astrophysical systems that operate on time-scales from seconds to years, data assimilation will be of fundamental importance when quantitative agreement between the model and observations is to be assessed. The ultimate verification that a model is correct is its ability to make reliable predictions, and for this data assimilation is necessary.

ACKNOWLEDGEMENTS

We acknowledge the NORDITA data assimilation program of 2011 for providing a stimulating scientific atmosphere. This work was supported by the European Research Council under the AstroDyn Research Project 227952. We would also like to thank the anonymous referee, Professor Spiegel and Dr Dobricic for their help in improving this paper.

REFERENCES

Bannister R., 2008a, *Q. J. R. Meteorol. Soc.*, 134, 1951
 Bannister R., 2008b, *Q. J. R. Meteorol. Soc.*, 134, 1971
 Bengtsson L., Ghil M., Källén E., 1981, *Dynamic Meteorology: Data Assimilation Methods*, Vol. 36. Springer, New York
 Brandenburg A., Dobler W., 2002, *Comput. Phys. Commun.*, 147, 471
 Brandenburg A., Jennings R. L., Nordlund Å., Rieutord M., Stein R. F., Tuominen I., 1996, *J. Fluid Mech.*, 306, 325

Brandenburg A., Chan K. L., Nordlund Å., Stein R. F., 2005, *Astron. Nachr.*, 326, 681
 Brun A. S., 2007, *Astron. Nachr.*, 328, 329
 Choudhuri A. R., Chatterjee P., Jiang J., 2007, *Phys. Rev. Lett.*, 98, 131103
 Cipra B., 1993, *SIAM News*, 26, 757
 Daley R., 1993, *Atmospheric Data Analysis*. Cambridge Univ. Press, Cambridge
 Dikpati M., 2007, 9th Annual Meeting of the Northwest Section of the APS, abstract E4.002
 Dintrans B., Brandenburg A., Nordlund Å., Stein R. F., 2005, *A&A*, 438, 365
 Dobricic S., Pinardi N., 2008, *Ocean Modelling*, 22, 89
 Ghil M., 1989, *Dyn. Atmos. Oceans*, 13, 171
 Hurlburt N. E., Toomre J., Massaguer J. M., 1986, *ApJ*, 311, 563
 Ilonidis S., Zhao J., Kosovichev A., 2011, *Sci*, 333, 993
 Jouve L., Brun A. S., Talagrand O., 2011, *ApJ*, 735, 31
 Judd K., Smith L., 2001, *Physica D: Nonlinear Phenomena*, 151, 125
 Judd K., Reynolds C., Rosmond T., Smith L., 2008, *J. Atmos. Sci.*, 65, 1749
 Kalman R., 1960, *J. Basic Eng.*, 82, 35
 Kalnay E., 2003, *Atmospheric Modelling, Data Assimilation, and Predictability*. Cambridge Univ. Press, Cambridge
 Kitiashvili I., Kosovichev A., 2008, *ApJ*, 688, L49
 Le Dimet F.-X., Talagrand O., 1986, *Tellus A*, 38, 97
 Lorenc A. C., 1981, *Mon. Weather Rev.*, 109, 701
 Lorenc A., 1986, *Q. J. R. Meteorol. Soc.*, 112, 1177
 Lorenz E. N., 1993, *The Essence of Chaos*. Univ. Washington Press, Seattle
 Palmer T., Hagedorn R., 2006, *Predictability of Weather and Climate*. Cambridge Univ. Press, Cambridge
 Pires C., Vautard R., Talagrand O., 1996, *Tellus A*, 48, 96
 Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes in FORTRAN. The Art Of Scientific Computing*. Cambridge Univ. Press, Cambridge
 Švanda M., Gizon L., Hanasoge S. M., Ustyugov S. D., 2011, *A&A*, 530, A148
 Talagrand O., 1997, *J. Meteorol. Soc. Japan*, 75, 191
 Wang X., Snyder C., Hamill T., 2007, *Mon. Weather Rev.*, 135, 222
 Warnecke J., Brandenburg A., Mitra D., 2011, *A&A*, 534, A11
 Wunsch C., 2006, *Discrete Inverse and State Estimation Problems: With Geophysical Fluid Applications*. Cambridge Univ. Press, Cambridge

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.