

Properties of the Ensemble Kalman Filter

Laura Baker

October 5, 2007

Abstract

Data assimilation is used in numerical weather prediction to improve weather forecasts by incorporating observation data into the model forecast. The Ensemble Kalman Filter (EnKF) is a method of data assimilation which updates an ensemble of states and uses this to provide a state estimate and associated error estimate at each step. Some implementations of the EnKF, such as the Ensemble Transform Kalman Filter (ETKF), have been found to exhibit undesirable behaviour. This dissertation investigates two such problems exhibited by the ETKF: ensemble collapse and bias. A solution to the ensemble collapse problem using random rotations is given and experimental results with the ETKF with the random rotations applied are presented. The conditions for an EnKF to be unbiased are examined, and these conditions are used to determine further restrictions on the random rotations to ensure that applying them to the filter will not introduce a bias. An unbiased filter

Declaration

I confirm that this is my own work, and the use of all material from other sources has been properly and fully acknowledged.

Laura Baker

Acknowledgments

I would like to thank my supervisor Sarah Dance for being about as helpful and enthusiastic as it is possible for a supervisor to be (and also for laughing

Contents

3.3.3	The EnKF algorithm	28
3.4	The Analysis Step of an Ensemble Square Root Filter	28
3.5	The Ensemble Transform Kalman Filter	29
3.6	Implementing the ETKF	30
4	Ensemble collapse	33
4.1	Previous work	33
4.2	Mathematical explanation of ensemble collapse	34
4.3	Random rotations	36
4.4	Skewness	37
5	Biased and unbiased filters	44
5.1	The bias problem	44
5.2	The Revised ETKF	46
5.3	Measuring bias	47
6	An unbiased filter with random rotations	53
6.1	A random rotation matrix which does not introduce bias	53
6.2	Experimental results	56
7	Conclusions	62
7.1	Summary	62
7.2	Future work	64
7.2.1	Initialisation experiments	64
7.2.2	A possible connection between ensemble collapse and bias	64
7.2.3	Optimal random rotations	65

List of Figures

1.1	Simplified diagram showing ensemble collapse	7
1.2	Simplified diagram showing a biased ensemble update	7
2.1	The swinging spring setup	11
2.2	Elastic and gravitational forces acting on the bob	11

6.2	ETKF with the new rotation matrix, ensemble skewness . . .	59
6.3	Revised ETKF with the new rotation matrix, ensemble statistics	60

List of Tables

2.1	Standard deviations used in the observation error covariance matrix and in generating the initial ensemble.	19
5.1	Average absolute value of bias for different filter formulations .	52
6.1	Average absolute value of bias for the ETKF and revised ETKF with the new random rotations	61

List of Abbreviations

DA	Data assimilation
EnKF	Ensemble Kalman Filter
ETKF	Ensemble Transform Kalman Filter
KF	Kalman Filter
NWP	Numerical Weather Prediction
SRF	Ensemble Square Root Filter
SVD	Singular Value Decomposition

Chapter 1

Introduction

1.1 Background

Data assimilation (DA) is a method of combining observation data with model forecast data in order to more accurately predict the state of a system. One of its most common uses is in numerical weather prediction (NWP). In NWP, we have observation data about different atmospheric properties, such as pressure, temperature and wind speed, obtained from various sources, such as ground stations, radiosondes, aircraft and satellites, at many different locations. We also have details of the state of the system at an earlier time, and a model which uses this to predict a forecast state. In the analysis step of the DA system, this model forecast state is combined with recent observation data to give an accurate estimate of the state of the system, called the analysis state. This analysis state is then used to produce the forecast state for the next analysis timestep.

The Kalman Filter (KF) is a sequential data assimilation scheme for use with linear systems. Each update step consists of a forecast step, in which the analysis from the previous step is run through the model, and an analysis step, in which the observation data is assimilated into the forecast from the model to produce the analysis. As well as updating the state estimate at each step, the KF also updates the error covariance, which is a measure of

the uncertainty in the state estimate.

The Ensemble Kalman Filter (EnKF), first described by Evensen [1994], generalises the idea of the KF to nonlinear systems by using an ensemble, or statistical sample, of forecasts. In the forecast step, the ensemble members are each forecast individually by the full nonlinear model, and in the analysis




Figure 1.1: Simplified diagram showing ensemble collapse

bias. Here the state estimate, which is updated separately from the individual ensemble members, is clearly not equal to the mean of the ensemble members.




Figure 1.2: Simplified diagram showing a biased ensemble update

1.3 Aims of this dissertation

The aim of this dissertation is to investigate the ensemble collapse and bias problems, and to try to develop a filter which does not exhibit either problem. This will be achieved by running experiments using the swinging spring system, which is described fully in Section 2.1, and analysing the results from

these experiments.

1.4 Principal new results

The main new results in this dissertation are

- Applying random rotations to the ETKF to eliminate the ensemble collapse problem
- Establishing that random rotations can introduce a bias to the filter
- Developing a new unbiased filter with random rotations

1.5 Outline of the dissertation

Chapter 2 introduces the swinging spring system, which is used in the experiments. Certain properties of this system make it relevant to NWP, and in particular to the problem of initialisation. The chapter describes the initialisation problem and investigates the initialisation of the swinging spring system. Details of the experimental setup are also given.

Chapter 3 gives a description of the EnKF and a particular implementation of it, the ensemble transform Kalman filter (ETKF). Details of how the filter is implemented in the experiments are also given.

In Chapter 4, we investigate the ensemble collapse problem. We begin by explaining why the problem occurs, and then give a solution to the problem, that of applying random rotations to the analysis update, which was suggested by Leeuwenburgh et al. [2005]. We also introduce ensemble skewness as a measure of ensemble collapse. We compare experimental results from the original ETKF and the ETKF with the random rotations applied.

Chapter 5 begins by describing the bias problem and why it occurs. The revised ETKF of Wang et al. [2004] is introduced, and experimental results comparing the ETKF and revised ETKF, both with and without the random

rotations of Chapter 4, are given. A quantitative measure of bias is also defined.

In Chapter 6 we develop an unbiased filter with random rotations by placing further restrictions on the random rotation matrix. The new filter is tested and compared with other filter formulations.

Finally, Chapter 7 gives a summary of the results and proposes some ideas for further work.

Chapter 2

Experimental setup and initialisation

The system used in the experiments is the swinging spring, or elastic pendulum system, which is described in Lynch [2003]. This system has both high and low frequency oscillations, and is therefore a good analogy of a model of atmospheric dynamics, in which both high frequency gravity waves and low frequency Rossby waves occur. It is relevant to the problem of initialisation in numerical weather prediction (NWP), which is described in Section 2.3. We begin by introducing the swinging spring system.

2.1 The swinging spring system

The experimental setup is shown in Figure 2.1. A heavy bob of mass m is suspended by a light elastic spring of elasticity k and unstretched length l_0 from a fixed point, and is allowed to move in a vertical plane. The motion of the bob is defined by the high frequency elastic oscillations of the spring, and lower frequency rotation oscillations about the suspension point.

The position of the bob at a given time is expressed in polar coordinates (r, θ) , where r is the distance from the suspension point and θ is the angle from the downward vertical direction. The radial and angular momenta are

given by $p_r = m\dot{r}$ and $p_\theta = mr^2\dot{\theta}$ respectively. The gravitational and elastic forces on the bob are as shown in Figure 2.2, where g is the gravitational acceleration.



Figure 2.1: The swinging spring setup. The bob is attached to an elastic spring, and its position is given by polar coordinates (r, θ) .




Figure 2.2: Elastic and gravitational forces acting on the bob, which has mass m .

We define the state vector $X = (\theta, p_\theta, r, p_r)^T$, and note that we can write this as $X = \begin{bmatrix} \Theta \\ R \end{bmatrix}$, where $\Theta = \begin{bmatrix} \theta \\ p_\theta \end{bmatrix}$ is the rotational component and $R = \begin{bmatrix} r \\ p_r \end{bmatrix}$

$$\begin{aligned}\dot{p}_r &= -\frac{\partial H}{\partial r} \quad , \quad \dot{r} = \frac{\partial H}{\partial p_r} \\ \dot{p}_\theta &= -\frac{\partial H}{\partial \theta} \quad , \quad \dot{\theta} = -\frac{\partial H}{\partial p_\theta}\end{aligned}$$

the following equations of motion for the bob can be derived:

$$\dot{\theta} = \frac{p_\theta}{mr^2} \quad (2.1)$$

$$\dot{p}_\theta = -mgr \sin \theta \quad (2.2)$$

$$\dot{r} = \frac{p_r}{m} \quad (2.3)$$

$$\dot{p}_r = \frac{p_\theta^2}{mr^3} - k(r - l_0) + mg \cos \theta. \quad (2.4)$$

The system has one stable equilibrium point, at which the bob is vertically below the suspension point and the spring is fully stretched. At this point the gravitational and elastic forces are balanced: $k(l - l_0) = mg$, where l is

the length of the spring at equilibrium. The equation of motion for the bob is

which has frequency $\omega_\theta = \sqrt{\frac{g}{l}}$, and the elastic component as

$$\ddot{r} + \frac{k}{m}r = 0,$$

which has frequency $\omega_r = \sqrt{\frac{k}{m}}$.

Defining η to be the ratio of the frequencies, we have

$$\eta = \frac{\omega_\theta}{\omega_r} = \sqrt{\frac{mg}{kl}} = \sqrt{\frac{l - l_0}{l}} < 1.$$

The parameter values for the experiments are chosen such that $\eta \ll 1$, i.e. the frequency of the rotational motion is much lower than that of the elastic motion.

2.2 Parameter values

For the purposes of our experiments, we use the parameter values $m = 1$, $g = \pi^2$, $k = 100\pi^2$ and $l = 1$, which are taken from the example in Lynch [2003]. This gives frequencies $\omega_\theta = \pi$ and $\omega_r = 10\pi$, so the frequency ratio is $\eta = 0.1$ and the cyclic frequencies of the rotational and elastic motions are $f_\theta = \omega_\theta/2\pi = 0.5$ and $f_r = \omega_r/2\pi = 5$ respectively. Figure 2.3 was produced using the numerical method detailed in Section 2.4. It shows the trajectories of each of the 4 variables plotted against time and the corresponding modulus of the Fourier transform coefficients, for the initial conditions $(\theta, p_\theta, r, p_r) = (1, 0, 1.05, 0)$. The Fourier transforms were computed using the MATLAB fast Fourier transform function `fft`. We can see that the θ and p_θ coordinates have large peaks in the Fourier transform plots at $f = f_\theta$ and the r and p_r coordinates have large peaks at $f = f_r$, but with small peaks at around $f = 2f_\theta$, indicating that there is some interaction between the fast and slow oscillations.

2.3 Initialisation

In the atmosphere there are many different motions with different timescales. In NWP systems, we are generally interested only in motions with timescale greater than one day, and motions with significantly smaller timescales are regarded as noise. Initialisation is the process of adjusting the initial condi-

amplitudes of the high frequency components to zero:

$$\mathbf{R} = \begin{pmatrix} r \\ p_r \end{pmatrix} = \mathbf{0}$$

at $t = 0$. So our initial conditions are $r(0) = l$ and $p_r(0) = 0$. This is called linear initialisation.

Figure 2.4 shows the trajectories and corresponding Fourier transforms for the swinging spring system with linear initialisation. It should be noted that the scales on the axes for the r and p_r variables are smaller than those for the uninitialised system. Compared to the uninitialised system, the θ and p_θ plots are unchanged, but the effect of the low frequency oscillations is having more of an effect on the frequency of the r and p_r coordinates, seen both by the slower oscillation in the trajectories as well as the high frequency oscillations, and by the larger peak at $f = 2f_\theta$ in the Fourier transform.

The linear initialisation has not had the desired result of removing the high frequency oscillations since the amplitude of the oscillations is too large for the linear system to be a good approximation to the system dynamics.

2.3.2 Nonlinear initialisation

9eJN5gyy,lesitf (9eJN5gyy,)1motions

$p_r(0) = 0$. For $\dot{p}_r = 0$ we substitute $\dot{\theta}(0)$, obtained from equation 2.1, into equation 2.4, set the right hand side of equation 2.4 equal to zero and solve for $r(0)$, giving

$$r(0) = \frac{l(1 - \frac{1}{2}(1 - \cos \theta(0)))}{1 - \frac{\dot{\theta}(0)^2}{\omega_r^2}}. \quad (2.9)$$

Finally, by equation 2.1 we have

$$p_\theta(0) = mr(0)^2\dot{\theta}(0). \quad (2.10)$$

Figure 2.5 shows the trajectories and corresponding Fourier transforms for the swinging spring system with nonlinear initialisation. As in the linear initialisation case, note that the scales on the axes for the r and p_r variables are again smaller. Again the θ and p_θ plots are unchanged, but now we can see that the high frequency oscillations in the r and p_r coordinates are almost no longer present, although very low amplitude high frequency oscillations can still be seen in the p_r coordinate, both by the small wiggles in the trajectory plot, and by the small peak in the Fourier transform at $f = f_r$. The

To determine the stability of the numerical method, we consider the stability of the linearised system given by equations 2.5-2.8. This should give

same as analysing the full nonlinear system, the parameter values chosen as a result of the linear stability analysis showed no instability in the experiments, both in this dissertation and in Livings [2005].

Although we are modelling a continuous system, for the purposes of our experiments we take the truth to be that generated by the discrete numerical approximation to the system, and not the truth from the continuous system which the equations are modelling. In the experiments, the discrete system is used both to generate the truth and in the forecast step.

2.5 Notes about the experimental setup

The experiments were carried out using the MATLAB code written for Livings [2005], though modifications and additions have been made for this dissertation.

In our experiments, we use frequent, perfect observations of all four coordinates. We take observations at time intervals of 0.1, beginning at time 0.1. Although the errors in the observations are actually zero, we assume standard deviations for the coordinates of around one tenth of the amplitude of the oscillations in the truth, and assume that the errors are uncorrelated. The standard deviations used are given in Table 2.1, and these determine the diagonal entries in the covariance matrix R . This covariance matrix is also used to generate the initial ensemble, which is formed by taking a pseudo-random sample of vectors from a normal distribution with mean equal to the true initial state and covariance matrix equal to this covariance matrix. The ensemble is then translated so that the ensemble mean is exactly equal to the true initial state. This last step represents having perfect observations at the initial time.

In the experiments in Chapters 4 and 5, the true state at time $t = 0$, and therefore also the ensemble mean at this time, are initialised using nonlinear initialisation; however, the ensemble at $t = 0$ is left uninitialised.

All the results presented in this dissertation use perfect observations and

a perfect model, and an ensemble of size $N = 10$.

Component	Standard deviation
θ	0.1
p_θ	0.3
r	




Figure 2.3: Trajectories for the swinging spring system and the corresponding Fourier transforms, with parameter values $m = 1$, $g = \pi^2$, $k = 100\pi^2$ and $l = 1$, and initial condition $Td[(\text{and})]TJ/F2811.5i3326(\text{condit})111.8f/42.5571.955Tf8.1380-32T8.7$




Figure 2.4: Trajectories for the swinging spring system with linear initialisation and the corresponding Fourier transforms, with parameter values $m = 1$, $g = \pi^2$, $k = 100\pi^2$ and $l = 1$, and initial conditions $(\theta, p_\theta, r, p_r) = (1, 0, 1.05, 0)$.




Figure 2.5: Trajectories for the swinging spring system with nonlinear initialisation and the corresponding Fourier transforms, with parameter values $m = 1$, $g = \pi^2$, $k = 100\pi^2$ and $l = 1$, and initial conditions $(\theta, p_\theta, r, p_r) = (1, 0, 1.05, 0)$.

Chapter 3

Description of the filter

In this chapter we describe the algorithms for the Kalman Filter (KF) and a general Ensemble Kalman Filter (EnKF), and describe a specific implementation of the EnKF, the Ensemble Transform Kalman Filter (ETKF), which is the implementation used for the experiments in this dissertation.

3.1 General Setup and Notation

Let n be the dimension of the state space, and let m be the dimension of the observation space. The state vector at a given time t_k is denoted by $\mathbf{x}(t_k)$, and the observation vector at the same time by $\mathbf{y}(t_k)$. The superscripts t , f and a are used to denote the truth, forecast and analysis respectively.

We assume that the system dynamics satisfy the model equation

$$\mathbf{x}^t(t_k) = \mathbf{m}_k(\mathbf{x}^t(t_{k-1})) + \boldsymbol{\eta}(t_{k-1}) \quad (3.1)$$

where $\mathbf{m}_k$ is the model operator, which can be linear or nonlinear, and $\boldsymbol{\eta}$ is the random model error, which is unbiased, and is from the normal distribution $N(0, \mathbf{Q}_k)$, where $\mathbf{Q}_k$ is a known covariance matrix. The observations satisfy

$$\mathbf{y}(t_k) = \mathbf{h}_k(\mathbf{x}^t(t_k)) + \boldsymbol{\varepsilon}(t_k) \quad (3.2)$$

where $\mathbf{h}_k$ is the observation operator, which again can be linear or nonlinear, and

3.2.1 The forecast step

The update equations for the state forecast and the error covariance forecast at time t_k are

$$\mathbf{x}_k^f = \mathbf{M}_k \mathbf{x}_{k-1}^a \quad (3.7)$$

$$\mathbf{P}_k^f = \mathbf{M}_k \mathbf{P}_{k-1}^a \mathbf{M}_k^T + \mathbf{Q}_k. \quad (3.8)$$

If we assume a perfect model then we can neglect the matrix $\mathbf{Q}_k$, and equation 3.8 becomes

$$\mathbf{P}_k^f = \mathbf{M}_k \mathbf{P}_{k-1}^a \mathbf{M}_k^T. \quad (3.9)$$

3.2.2 The analysis step

The update equations for the analysis state and error covariance are

$$\mathbf{x}_k^a = \mathbf{x}_k^f + \mathbf{K}_k [\mathbf{y}_k - \mathbf{H}_k \mathbf{x}_k^f] \quad (3.10)$$

$$\mathbf{P}_k^a = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k^f \quad (3.11)$$

where $\mathbf{H}_k$ is the linear observation operator and $\mathbf{K}_k$ is the Kalman gain

3.3.1 Notation

To describe the EnKF, we first introduce some new notation.

We use an ensemble of size N , and denote the ensemble members by $\mathbf{x}_i$, $i = 1, \dots, N$, each of which is of length n . The ensemble mean is given by

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i. \quad (3.14)$$

We define the $n \times N$ ensemble matrix by

$$\mathbf{X} = \sqrt{\frac{1}{N-1}} \begin{bmatrix} \mathbf{x}_1 & \dots & \mathbf{x}_N \end{bmatrix} \quad (3.15)$$

and the ensemble perturbation matrix by

$$\mathbf{X}^\perp = \sqrt{\frac{1}{N-1}} \begin{bmatrix} \mathbf{x}_1 - \bar{\mathbf{x}} & \dots & \mathbf{x}_N - \bar{\mathbf{x}} \end{bmatrix}. \quad (3.16)$$

The ensemble covariance matrix is of dimension $n \times n$ and is given by

$$\mathbf{P}_e = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T$$

We define an ensemble perturbation matrix $Y^{\square}$, with columns $y_i^{\square}$, $i = 1, \dots, N$, which can be written as

$$Y^{\square} = H X^{\square}. \quad (3.21)$$

Using this equation, we can write the update equations 3.10 and 3.11 without including the observation operator matrix H , and since both $Y^{\square}$ and $X^{\square}$ are perturbation matrices, the equations will then be valid for both linear and nonlinear observation operators. (Livings et al. [2006])

We define an ensemble version of the Kalman gain matrix by

$$K_e = P_e^f H^T \square H P_e^f H^T + R^{\square -1}. \quad (3.22)$$

Using equations 3.21 and the fact that $P^f = X^{\square f} X^{\square f T}$, we can rewrite this in terms of $X^{\square f}$ and $Y^{\square f}$ as

$$\begin{aligned} K_e &= P_e^f H^T \square H P_e^f H^T + R^{\square -1} \\ &= X^{\square f} X^{\square f T} H^T \square H X^{\square f} X^{\square f T} H^T + R^{\square -1} \\ &= X^{\square f} Y^{\square f T} \square Y^{\square f} Y^{\square f T} + R^{\square -1} \\ &= X^{\square f} Y^{\square f T} S^{-1} \end{aligned} \quad (3.23)$$

where

$$S = Y^{\square f} Y^{\square f T} + R. \quad (3.24)$$

Similarly, we can write the covariance update equation 3.11 as

$$\begin{aligned} P_e^a &= (I - K_e H) P_e^f \\ &= \square I - X^{\square f} Y^{\square f T} S^{-1} H \square X^{\square f} X^{\square f T} \\ &= X^{\square f} \square I - Y^{\square f T} S^{-1} H X^{\square f} \square X^{\square f T} \\ &= X^{\square f} \square I - Y^{\square f T} S^{-1} Y^{\square f} \square X^{\square f T}. \end{aligned} \quad (3.25)$$

An alternative method of applying the EnKF to the case of nonlinear observation operators is given in Section 4.5 of Evensen [2003].

3.3.3 The EnKF algorithm

The main steps of a general EnKF algorithm are as follows:

- Given an initial state estimate $\mathbf{x}_0^a$ and error covariance matrix $\mathbf{P}_0^a$, we begin by generating an analysis ensemble of initial states $\mathbf{x}_i^a$, $i = 1, \dots, N$, whose mean is equal to $\mathbf{x}_0^a$ and whose covariance is determined by $\mathbf{P}_0^a$.
- The forecast step: the ensemble members are each updated using the full nonlinear dynamical model:

$$\mathbf{x}_i^f = \mathbf{m}_k(\mathbf{x}_i^a) \quad (3.26)$$

and the mean $\bar{\mathbf{x}}^f$ and covariance $\mathbf{P}_e^f$ of the forecast ensemble are computed.

- The analysis step: the mean and covariance of this forecast ensemble are then used to assimilate the observations and the result is used to compute a new analysis ensemble.

There are many different formulations of the EnKF, and these differ in the analysis step. One method, the perturbed observation EnKF described in Burgers et al. [1998], generates an ensemble of observations consistent with the error statistics of the observation and assimilates these into each ensemble member. Many other formulations come under the category of ensemble square root filters (SRFs). Tippett et al. [2003] gives a uniform framework for SRFs.

3.4 The Analysis Step of an Ensemble Square Root Filter

In Section 3.3.2 we defined an ensemble version of the Kalman gain matrix and wrote this as

$$\mathbf{K}_e = \mathbf{X}^{\parallel f} \mathbf{Y}^{\parallel f T} \mathbf{S}^{-1}. \quad (3.27)$$

The analysis state estimate is updated by

$$\tilde{\mathbf{x}} =$$

computationally expensive inversion of the matrix S , the ETKF makes use of the fact that the matrix R is much easier to invert since it has a simpler structure, normally diagonal.

We use the identity

$$I - Y^{\text{of}T} S^{-1} Y^{\text{of}} = [I + Y^{\text{of}T} R^{-1} Y^{\text{of}}]^{-1}, \quad (3.33)$$

which can be shown by postmultiplying both sides of the equation by $I + Y^{\text{of}T} R^{-1} Y^{\text{of}}$ and substituting in equation 3.24. Using this, we can avoid the problem of computing S^{-1} .

We compute the eigenvalue decomposition

$$Y^{\text{of}T} R^{-1} Y^{\text{of}} = C \Lambda C^T \quad (3.34)$$

where C is a matrix of orthonormal eigenvectors and Λ is the diagonal matrix of corresponding eigenvalues. Then

$$\begin{aligned} I - Y^{\text{of}T} S^{-1} Y^{\text{of}} &= [I + C \Lambda C^T]^{-1} \\ &= C (I + \Lambda)^{-1} C^T \end{aligned}$$

and so we can take

$$T = C (I + \Lambda)^{-\frac{1}{2}} \quad (3.35)$$

as the desired matrix square root. Thus the analysis ensemble perturbation matrix is updated by

$$X^{\text{a}} = X^{\text{f}} C (I + \Lambda)^{-\frac{1}{2}}. \quad (3.36)$$

Since Λ

To avoid storing the Kalman gain matrix K_e , we build up the product vector

d

Chapter 4

Ensemble collapse

In this chapter we discuss the problem of ensemble collapse, which some formulations of EnKFs, such as the ETKF, have been found to exhibit. We begin by reviewing results from previous work on this subject in Section 4.1. In Section 4.2 we give a mathematical explanation of why the ensemble collapse problem occurs, and Section 4.3 examines one proposed solution to the problem, that of applying random rotations, and gives results of experiments with this solution applied. Finally, Section 4.4 suggests the use of skewness as a quantitative measure of ensemble collapse and gives some experimental results.

4.1 Previous work

Livings [2005] carried out experiments with the swinging spring system described in Section 2.1 using the same experimental setup as in this dissertation, which is detailed in Chapter 2. The experiments used the ETKF formulation of the EnKF, which is described in Section 3.5 and is the filter formulation used in the experiments for this dissertation. The results, assuming a perfect model, with frequent observations of all coordinates, and an ensemble of size $N = 10$, showed that there was a collapse from 10 ensemble members to just 5. Figure 4.1 is a recreation of the results from Livings

[2005], produced using the same MATLAB code. It shows the individual ensemble members relative to the truth for each of the 4 coordinates, and it can be seen that there is a decrease after the first assimilation from 10 distinct ensemble members to 5.

Lawson and Hansen [2004] performed experiments with the Ikeda system, a two-dimensional, nonlinear, chaotic system, which is described in more detail in Section 3b of Lawson and Hansen [2004]. The filter used was the EnSRF, a square root filter which is analytically equivalent to the ETKF of

problem occurs, but here we follow the more rigorous explanation of Livings [2005].

The following theorem explains why ensemble collapse occurs in the ETKF.

Theorem 1 *For a system with $N > m$, the number of distinct ensemble members is reduced to at most $m + 1$, with at least $N - m$ ensemble members equal to the ensemble mean $\bar{x}^a$ (i.e. the perturbation from the mean is zero) and the remaining ensemble members having different states.*

Proof This proof is based on Section 6.2 of Livings [2005]. Suppose $N > m$ and let the observation operator H be linear. Then writing $Y^{\text{la}} = H X^{\text{la}}$ we can rewrite the update equation for the analysis ensemble perturbation as

$$\begin{aligned} Y^{\text{la}} &= H X^{\text{lf}} C (I + \Lambda)^{-\frac{1}{2}} \\ &= Y^{\text{lf}} C (I + \Lambda)^{-\frac{1}{2}} \\ &= R^{\frac{1}{2}} \Psi^f C (I + \Lambda)^{-\frac{1}{2}} \\ &= R^{\frac{1}{2}} V \Sigma^T (I + \Lambda)^{-\frac{1}{2}} \end{aligned}$$

where C , Σ and V are from the SVD 3.40. Σ is an $N \times m$ matrix of the form

$$\Sigma = \begin{bmatrix} & \\ & \Sigma_1 \\ & \\ 0 & \end{bmatrix} \quad (4.3)$$

where Σ_1 is an $m \times m$ diagonal matrix and 0 denotes the $(N - m) \times m$ submatrix of zeros. The $N \times N$ matrix Λ is defined by

$$\Lambda = \Sigma \Sigma^T = \begin{bmatrix} & \\ & \Lambda_1 & 0 \\ & 0 & 0 \end{bmatrix} \quad (4.4)$$

where $\Lambda_1 = \Sigma_1 \Sigma_1^T$. Thus the number of non-zero entries on the diagonals of Σ_1 and Λ_1 correspond to the number of non-zero eigenvalues of Ψ^{fT} .

Substituting into the expression for Y^{la} gives

$$\begin{aligned} Y^{\text{la}} &= R^{\frac{1}{2}} V \begin{bmatrix} & \\ & \Sigma_1^T (I + \Lambda_1)^{-\frac{1}{2}} & 0 \\ & \\ 0 & \end{bmatrix} \\ &= R^{\frac{1}{2}} V \Sigma_1^T (I + \Lambda_1)^{-\frac{1}{2}} \begin{bmatrix} & \\ & 0 \\ & \\ 0 & \end{bmatrix} \end{aligned} \quad (4.5)$$

where here $\mathbf{0}$ denotes the $m \times (N - m)$ matrix of zeros and $\mathbf{I}$ is the $m \times m$ identity matrix. Thus $\mathbf{Y}^{\text{la}}$ can have no more than m non-zero columns, which is equivalent to saying that no more than m ensemble members can differ from $\bar{\mathbf{x}}^a$, and therefore the ensemble members can have no more than $m + 1$ distinct states. $\square$

Figure 4.2 shows the ensemble members relative to the truth, plotted at each of the first 10 assimilations. This shows the ensemble collapse property of the filter in more detail than Figure 4.1. We can see that the number of distinct ensemble members reduces to 5 after the first assimilation, and decreases further after subsequent assimilations.

It should be noted that in NWP systems, we typically have a very large observation space, and so $N \gg m$. Therefore the ensemble collapse problem will not affect such systems. However, in some cases the operations are processed in small batches (Brown and Hwang [1997]) and so the value of m at a given time may be small. In such systems, and in systems with low dimensional observation spaces such as the swinging spring system, it should be considered as a potential problem if using an ensemble size $N > m + 1$.

4.3 Random rotations

The random rotations suggested by Leeuwenburgh et al. [2005] are achieved by postmultiplying equation 3.36 by an $N \times N$ random orthogonal matrix $\mathbf{U}$. If we define an $m \times m$ matrix $\mathbf{A}$ by

$$\mathbf{A} = \mathbf{R}^{\frac{1}{2}} \mathbf{V} \Sigma_1^T (\mathbf{I} + \Lambda_1)^{-\frac{1}{2}} \quad (4.6)$$

then we can write equation 4.5 as

$$\mathbf{Y}^{\text{la}} = \begin{bmatrix} \mathbf{A} & \mathbf{0} \end{bmatrix}. \quad (4.7)$$

We can write the matrix $\mathbf{U}$ in terms of submatrices as

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \\ \mathbf{U}_3 & \mathbf{U}_4 \end{bmatrix} \quad (4.8)$$

where U_1 , U_2 , U_3 and U_4 are of dimensions $m \times m$, $m \times (N - m)$, $(N - m) \times m$ and $(N - m) \times (N - m)$ respectively. Then

$$Y^{\text{la}}U = \begin{bmatrix} AU_1 & AU_2 \end{bmatrix}. \quad (4.9)$$

Since the matrix U can be any orthogonal matrix, this does not guarantee an increase in the number of non-zero columns of the matrix; for example, in the case $U = I$ the matrix will remain the same. However, in most cases there is an increase in the number of non-zero columns compared with the original matrix.

In this implementation, the matrix U is obtained from the SVD $U \Delta W = B$ where B is a pseudo-randomly generated matrix whose entries are from a uniform distribution on the interval $(0, 1)$. This is produced using the MATLAB function `rand(N)`.

Figure 4.3 shows the results from the ETKF with random rotations. The ensemble members are plotted relative to the truth, at each of the first 10 assimilations. The rotations appear to have stopped the ensemble collapse problem from occurring.

4.4 Skewness

Skewness is a measure of the asymmetry of a distribution. A perfectly symmetric, Gaussian distribution would have zero skewness. Since in the case of ensemble collapse we will have a non-Gaussian distribution of ensemble members, we can use skewness as a means of quantitatively measuring ensemble collapse.

For now we consider x_i to be one component of the ensemble member vector x_i . We use the definition of skewness given in Appendix A of Lawson and Hansen [2004]. For an ensemble $x_1, \dots, x_N$, we define the mean, variance and third central moment by

$$\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$$




Figure 4.1: ETKF, ensemble members. Coordinates are plotted relative to the truth. Red lines show ensemble members, black lines are observations plotted as error bars.




Figure 4.2: ETKF, ensemble members at each analysis step. Ensemble members are plotted as coloured stars relative to the truth at each of the first 10 assimilations.




Figure 4.3: ETKF with random rotations at each analysis step. Ensemble members are plotted as coloured stars relative to the truth at each of the first 10 assimilations. The ensemble collapse seen in the original ETKF is no longer occurring.




Figure 4.4: Absolute value of ensemble skewness for each coordinate plotted against time for the ETKF.




Figure 4.5: Absolute value of ensemble skewness for each coordinate plotted against time for the ETKF with random rotation.

Chapter 5

Biased and unbiased filters

Since X^{la} is a perturbation matrix, we should find that its column sum, given by $X^{\text{la}}\mathbf{1}$, is equal to zero, where $\mathbf{1} = (1, \dots, 1)^T$ of length N . However,

$$X^{\text{la}}\mathbf{1} = -\sqrt{\frac{1}{N-1}} \sum_{i=1}^N (x_i^a - \tilde{x})$$

Thus the covariance matrix is given by

$$\begin{aligned}
\mathbf{P}_e &= \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T \\
&= \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i^{\parallel} - \bar{\mathbf{x}}^{\parallel})(\mathbf{x}_i^{\parallel} - \bar{\mathbf{x}}^{\parallel})^T \\
&= \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i^{\parallel} \mathbf{x}_i^{\parallel T} - \mathbf{x}_i^{\parallel} \bar{\mathbf{x}}^{\parallel T} - \bar{\mathbf{x}}^{\parallel} \mathbf{x}_i^{\parallel T} + \bar{\mathbf{x}}^{\parallel} \bar{\mathbf{x}}^{\parallel T}) \\
&= \mathbf{X}^{\parallel a} \mathbf{X}^{\parallel a T} - \frac{N}{N-1} \bar{\mathbf{x}}^{\parallel} \bar{\mathbf{x}}^{\parallel T}.
\end{aligned}$$

If the filter is unbiased then $\bar{\mathbf{x}}^{\parallel} = \mathbf{0}$ and the last line is just

$$\mathbf{X}^{\parallel a} \mathbf{X}^{\parallel a T}.$$

Otherwise we have a smaller covariance than we expect, and therefore the standard deviation for each coordinate will also be too small. This results in the filter being overconfident. Figure 5.1 shows results from the ETKF. The ensemble mean and ensemble mean $\pm$ standard deviation are plotted relative to the truth. The truth is frequently more than one standard deviation away from the ensemble mean, which implies the filter is biased and overconfident.

5.2 The Revised ETKF

Wang et al. [2004] suggest a new filter, the revised ETKF, obtained modifying the ETKF by postmultiplying equation 3.36 by the transpose of the orthogonal matrix $\mathbf{C}$. The resulting analysis update equation is

$$\mathbf{X}^{\parallel a} = \mathbf{X}^{\parallel f} \mathbf{C} (\mathbf{I} + \mathbf{\Lambda})^{-\frac{1}{2}} \mathbf{C}^T. \quad (5.6)$$

Note that our new matrix $\mathbf{T} = \mathbf{C} (\mathbf{I} + \mathbf{\Lambda})^{-\frac{1}{2}} \mathbf{C}^T$ is symmetric. Theorem 2 of Livings et al. [2007] states that if $\mathbf{T}$ is symmetric then the filter with analysis update $\mathbf{X}^{\parallel a} = \mathbf{X}^{\parallel f} \mathbf{T}$ is unbiased, and therefore we can deduce that the revised ETKF is unbiased. See Livings et al. [2007] Section 3 for the proof of this theorem.

Figure 5.2 shows the ensemble mean and ensemble mean $\pm$ standard deviation plotted relative to the truth for the revised ETKF. The ensemble mean is now almost always within one standard deviation of the truth, so this filter is not overconfident.

Since the matrix $\mathbf{C}^T$ is orthogonal, the revised ETKF also overcomes the ensemble collapse problem, by the argument of Section 4.3. Figure 5.3, which shows the ensemble members plotted as coloured stars relative to the truth at each analysis time, and it can be seen that there is no evidence of ensemble collapse.

5.3 Measuring bias

In the same way that we used the ensemble skewness as a quantitative measure of ensemble collapse in Section 4.4, it is also useful to measure the biasedness of the filter in a similar way. Since an unbiased filter has the column sum of the perturbation matrix $\mathbf{X}^{\text{la}}$ equal to zero, we use this column sum to measure the biasedness of the filter. We define the bias b by

$$b = \frac{1}{N}$$

values for the other filter formulations are much larger, indicating that they




Figure 5.1: ETKF, ensemble statistics. Ensemble is mean shown in red, and ensemble mean $\pm$ standard deviation is shown in blue, both relative to the truth. The filter is biased and overconfident.





Figure 5.2: Revised ETKF, ensemble statistics. Ensemble mean is shown in red, and ensemble mean $\pm$ standard deviation is shown in blue, both relative to the truth. The signs of bias and overconfidence seen in Figure 5.1 are no longer present.



Filter formulation	Average $ b $ over 60 assimilations			
	θ	p_θ	r	p_r
ETKF	0.0045	0.0190	0.0001	0.0008
Revised ETKF	0.0452×10^{-14}	0.1517×10^{-14}	0.0273×10^{-14}	0.0025×10^{-14}
ETKF with random rotations	0.0053	0.0118	0.0001	0.0007
Revised ETKF with random rotations	0.0045	0.0190	0.0001	0.0008

Table 5.1: Average absolute value of bias for different filter formulations over 60 assimilations, for each of the 4 components.

Chapter 6

An unbiased filter with random rotations

by the fact that if we write λ as the corresponding eigenvalue, then

$$\mathbf{X}^a \mathbf{1} = \mathbf{X}^f \mathbf{T} \mathbf{1} = \lambda \mathbf{X}^f \mathbf{1} = \mathbf{0}. \quad (6.1)$$

So in order to find a rotation matrix which will not introduce a bias, we need to construct an $N \times N$ orthogonal matrix $\mathbf{U}$ such that $\mathbf{1} = (1, \dots, 1)^T$ is an eigenvector. To do this, we use the following theorem:

Theorem 2 Let $\mathbf{W}$ be an $N \times N$ orthogonal matrix whose first column is proportional to $\mathbf{1} = (1, \dots, 1)^T$

so $\mathbf{1}$ is an eigenvector of $\mathbf{U}$, as required. $\square$

So to construct $\mathbf{U}$ with the required properties, we first construct the matrices $\mathbf{V}$ and $\mathbf{W}$ such that $\mathbf{V}$ is of the form

$$\mathbf{V} = \begin{bmatrix} 1 & 0 \\ 0 & \mathbf{V}_1 \end{bmatrix}$$

and $\mathbf{W}$ has first column equal to $\alpha \mathbf{1}$; and then find $\mathbf{U}$ by solving the equation

$$\mathbf{U} = \mathbf{W} \mathbf{V} \mathbf{W}^T.$$

$\mathbf{V}_1$ is found from the SVD of a random $(N - 1) \times (N - 1)$ matrix $\mathbf{B}$:
 $\mathbf{B} = \mathbf{U}_1 \mathbf{\Sigma}_1 \mathbf{V}_1$.

$\mathbf{W}$ is constructed by applying the Gram-Schmidt Orthogonalisation algorithm to the matrix

Theorem 3 of Livings et al. [2007] states that if equation 3.30 gives an unbiased filter then equation 3.32 also gives an unbiased filter if the matrix U has $\mathbf{1}$ as an eigenvector. This is proved by the fact that if $X^{\text{rf}} T \mathbf{1} = \mathbf{0}$ then

$$X^{\text{rf}} T U \mathbf{1} = \nu X^{\text{rf}} T \mathbf{1} = \mathbf{0}. \quad (6.3)$$

where ν is the corresponding eigenvalue. Therefore if we apply the new random rotations to an unbiased filter such as the revised ETKF, we expect that the modified filter will still be unbiased.

6.2 Experimental results

Ideally we would like to test this postmultiplier matrix by applying it to an unbiased filter which has the ensemble collapse problem. In this case, we would expect that the modified filter would still be unbiased, and that it would no longer have the ensemble collapse problem. However, here we test the two properties separately, using the ETKF to test for ensemble collapse and skewness, and the revised ETKF to test for bias.

Figure 6.1 shows the results of the ETKF with the new rotation matrix applied. The ensemble members are plotted relative to the truth, at each of the first 10 assimilations, as in Figure 4.2. The modified filter does not show signs of ensemble collapse. The absolute value of the ensemble skewness for the modified filter is shown in Figure 6.2. The values of absolute skewness

applied. The values for the revised ETKF with the new rotations indicate that the filter does remain unbiased. However, the ETKF with the new rotations is still biased.




Figure 6.1: ETKF with the new rotation matrix. Ensemble members are plotted as coloured stars relative to the truth at each of the first 10 assimilations.




Figure 6.2: Absolute value of ensemble skewness for each coordinate plotted against time for the ETKF with the new rotation matrix.




Figure 6.3: ETKF with the new rotation matrix, ensemble statistics. Ensemble mean is shown in red, and ensemble mean $\pm$ standard deviation is shown in blue, both relative to the truth.

Filter formulation	Average $ b $ over 60 assimilations			
	θ	p_θ	r	p_r
ETKF with new random rotations	0.0045	0.0190	0.0001	0.0008
Revised ETKF with new random rotations	0.0400×10^{-14}	0.1963×10^{-14}	0.0278×10^{-14}	0.0023×10^{-14}

Table 6.1: Average absolute value of bias for the ETKF and revised ETKF with the new random rotations, over 60 assimilations, for each of the 4 components.

Chapter 7

Conclusions

7.1 Summary

In this dissertation, we have investigated the problems of ensemble collapse and bias which are seen in some implementations of the EnKF, including the ETKF. We have developed a modified version of the ETKF which does not exhibit either problem.

In Chapter 2, we introduced the swinging spring system, which can be used as an analogy of a model of atmospheric dynamics. In particular, it is relevant to the problem of initialisation in NWP. We described linear and nonlinear initialisation methods, and illustrated results for each of these for the swinging spring system.

Chapter 3 introduced the EnKF and the particular EnKF implementation, the ETKF, which was the filter used in the experiments in later chapters.

In Chapter 4, we investigated the problem of ensemble collapse. We

of the ensemble skewness for the two filters. The ETKF had larger ensemble skewness than the ETKF with random rotations. These results agreed with those of Leeuwenburgh et al. [2005].

In Chapter 5, we described the bias problem, and introduced the revised ETKF of Wang et al. [2004]. While the ETKF was seen to have a bias in the ensemble mean, the revised ETKF was not. The revised ETKF was also shown not to have the ensemble collapse problem. However, it was found that applying the random rotations described in Chapter 4 to the revised ETKF introduced a bias in the results.

In Chapter 6, using results from Livings et al. [2007], we developed further restrictions on the random rotation matrix to ensure that it would not introduce a bias when applied to an unbiased filter. The new rotation matrix was tested by applying it to both the ETKF and the revised ETKF. It was found that the ETKF with the new rotations did not exhibit the ensemble collapse problem seen in the original ETKF, and that applying the new rotations to the revised ETKF did not introduce a bias to the results. These results imply that in the case of an unbiased filter with the ensemble collapse problem, the new random rotations could be applied to prevent ensemble collapse without introducing a bias.

It should be noted that the experimental setup used in this dissertation was not ideal. Firstly, the experimental results shown here were all for just one run of the filter, with one particular initial state, set of observations and initial ensemble. It would be useful to carry out more runs of the filter to check that the results still agree with those presented here. Also the ensemble size used was always 10, so the experiments should be repeated with different ensemble sizes. The experiments here assumed a perfect model and used perfect observations of all 4 components at each assimilation time, but in a more realistic situation, this would not be the case. Therefore it would be useful to carry out the experiments on an imperfect system.

7.2 Future work

7.2.1 Initialisation experiments

We now have two filters which do not have either of the undesirable properties of bias or ensemble collapse: the original revised ETKF of Wang et al. [2004], described in Section 5.2, and the revised ETKF with the restricted random rotations of Section 6.1. Although the extra matrix multiplication of the revised ETKF with the restricted random rotations adds an additional computational cost to the filter compared to the revised ETKF, it is possible that there are other properties which are improved by the rotations, for example the initialisation properties of the two filters. It would be interesting to compare the two filters.

This could be done by repeating the assimilation up to a certain time, and then leaving the model to forecast for the remaining time with no further observations. The results could then be compared for the two filters to see if there is a difference in how well-initialised the system is after the assimilations, ie. whether the amplitude of the fast oscillations remain small.

but as well as being unbiased, the revised ETKF also no longer has the ensemble collapse problem. There may be a connection between the two properties, and it could in fact be true that an unbiased filter will never have the ensemble collapse problem, or conversely that a filter with the ensemble collapse problem must always be biased.

This possible connection could be investigated using the conditions for unbiasedness in Livings et al. [2007] and the explanation of ensemble collapse in Section 4.2.

7.2.3 Optimal random rotations

Another way of developing the random rotations idea further would be to try to find an optimal random rotation matrix. In Section 4.4 we used the ensemble skewness as a measure of ensemble collapse, so we could try to find a random rotation matrix such that the ensemble skewness is minimised.

Bibliography

J. L. Anderson. An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.*, 129:2884–2903, 2001.

C. H. Bishop, B. J. Etherton, and S. H. Majumdar. Adaptive sampling with the ensemble transform Kalman filter, part I: Theoretical aspects. *Mon. Wea. Rev.*, 129:420–436, 2001.

R. G. Brown and P. Y. C. Hwang. *Introduction to random signals and applied*

