

On the orthogonality of Bred Vectors

J.D. Annan¹

Frontier Research System for Global Change

3173-25 Showa-machi Kanazawa-ku

Yokohama, Kanagawa 236-0001, Japan

and

Proudman Oceanographic Laboratory

Joseph Proudman Building

6 Brownlow Street, Liverpool, L3 5DA, UK

`jdannan@jamstec.go.jp`

31 October 2003

¹Corresponding author

Abstract

The ‘Breeding Method’ is a well established and computationally cheap method for generating perturbations for ensemble integrations. One feature of bred vectors is that, although arbitrary initial perturbations generally converge towards the fastest growing directions, distinct bred vectors for global numerical weather prediction (NWP) models are generally quasi-orthogonal. However, an examination of the local structure of the vectors indicates that this conclusion may be somewhat misleading, since the perturbations frequently have very similar shapes over local areas (consistent with the low local dimension of atmospheric physics). Therefore, there may be substantial redundancy when multiple independent breeding cycles are performed in parallel, and the vectors can be inefficient in spanning the range of locally growing perturbations. Experiments with the 3 variable Lorenz63 model indicate that orthogonalising the bred vectors can result in significantly improved performance. A simple local orthogonalisation technique has been applied to the spatially resolved 40 variable Lorenz96 model with encouraging results. It is suggested that a similar treatment to bred vectors might be helpful for global NWP or other applications.

1 Introduction

The ‘Breeding method’ for generating perturbations for ensemble perturbations was introduced by Toth & Kalnay (1993). In this technique, alongside the main forecast, a perturbed model run is also integrated. The bred vector is defined to be the difference between the perturbed forecast and the control run. At regular intervals, the perturbation is rescaled to a specific magnitude (usually chosen to be comparable to the analysis error, and allowed to slowly vary over the model domain). For any initial perturbation, the bred vector rapidly turns towards the fastest growing perturbations of this size, and can be considered as a finite size equivalent of the leading Lyapunov vector (Kalnay *et al.*, 2002). The technique can be generalised to an arbitrary number of independent perturbed model runs. As well as simply predicting the reliability of forecasts, these perturbations have many related uses, including superior forecasting skill through ensemble averaging (Toth & Kalnay, 1997), and improvement of the quality of analyses through feed back into the prior covariance matrix (Kalnay & Toth, 1994; Corazza *et al.*, 2002). These bred vectors depend strongly on the ‘flow of the day’ and therefore substantially outperform random perturbations.

One superficially attractive property of the bred vectors is that although different initial perturbations all generally turn towards the fastest growing modes, they are globally quasi-orthogonal (Toth & Kalnay, 1997) and therefore (one might expect) automatically provide a good sample of the different dominant growing

error directions, without the need for additional computation. However, a look at the local structure in regions of strong error growth indicates that this global orthogonality is somewhat misleading. In fact, the perturbations are often locally rather similar in shape, differing only in sign and amplitude. For example, Table 1 in Toth & Kalnay (1997) shows that when 20 independent bred vectors were analysed across 6 subregions, comfortably more than half of the local perturbations are near-replicates of the most popular shapes (and presumably there is further duplication amongst the remaining perturbations). Since the sign of the local perturbation is arbitrary, any two independent bred vectors will tend to have locally strong positive and negative correlations which will approximately cancel out when averaged globally. It would appear from this that there may be substantial duplication of computing effort in their repeated calculation. Furthermore, as we shall see in subsequent sections, the locally narrow distribution of bred vectors can result in them temporarily failing to track the optimal perturbations as they evolve through time. The main goal of this paper is to investigate the extent to which the performance of bred vectors could be improved by ensuring that they are locally, and not just globally, orthogonal.

2 Orthogonal breeding in the Lorenz63 model

We start by using the famous 3-variable Lorenz63 model (Lorenz, 1963) to investigate the performance of orthogonalised versus standard bred perturbations. Here the Lorenz63 model can be thought of as a proxy for a small region of rapid error growth, rather than the entire global atmosphere. The basic parameters for our first experiments are a rescaling interval of 0.1 time units, and a rescaled perturbation size (in the L_2 norm) of 1. In order to prevent the bred vectors from converging completely (and to simulate the fine scale stochastic forcing of *eg* convection), small uncorrelated random perturbations of magnitude 0.01 are added to each variable. The qualitative results are not very sensitive to any of these choices.

For the orthogonalisation experiments, we use modified Gram-Schmidt orthogonalisation (Kincaid & Cheney, 1990) with (at least initially) the vectors ordered according to their size. This is the finite interval, finite size analogue to the method used by Yamane & Yoden (2001) for the calculation of the leading backward Lyapunov vectors. The treatment of the leading vector is therefore identical to the standard breeding method, ie it is rescaled with no change in direction. The second vector is projected onto the subspace orthogonal to the leading vector before being rescaled, and iteratively the n th vector is projected onto the subspace orthogonal to the previous $n - 1$ vectors before being rescaled.

The first point of comparison is simply the growth of the vectors from the two schemes, that is the ratio of final to initial lengths over the integration interval (for an initial length of 1, this is simply the final length of the vector). The standard breeding method results in an average ratio of 1.18, and clearly this will not vary with the number of vectors generated. When orthogonalisation is performed, the results are, of course, identical for a single vector, but for two vectors, the fastest growing of the pair has an average final length of 1.44, and the second vector shrinks to a length of 0.93. Adding a third orthogonal vector (which has an average final length of 0.44) has a negligible effect on the growth of the leading two vectors. The mean absolute correlation between the leading pair of orthogonal vectors grows from zero to around 0.4 at the end of the integration interval.

Perhaps a fairer comparison would be to look at the mean largest vector out of an ensemble of independently bred vectors. In this case, the instantaneously fastest growing bred vector from an ensemble size of 3 grows to an average length of 1.32. It is not until we have around 10 independently bred perturbations that the highest growth factor found matches that of the orthogonal pair. So from this simple test it appears that orthogonal vectors are substantially more efficient at generating fast growing perturbations. The reason for this is that the growth factors of the leading orthogonalised perturbations vary through time, and frequently the growth factor of the second perturbation overtakes that of the first one. At this

point, the orthogonalised system immediately captures the rapid growth of the new leading perturbation. However, in the case of standard bred vectors, the perturbations must rotate through anything up to 90° to align with the direction of the new fastest growing perturbation. The typical correlation between any two independently bred vectors in this experiment is around 0.8, so many of them are required in order for one of them to coincide with the new rapidly growing direction. During the time it takes the vectors to rotate, their growth rates may be substantially depressed relative to the fastest growing vector from the orthogonal pair.

This is illustrated in Figure 1 which shows a comparison of two independent bred vectors (BV1, BV2) with a pair of orthogonal bred vectors (OV1, OV2) through 2.5 time units during one of the experiments. The upper graph shows how the length of each vector varies through time, with all being rescaled to unit length at intervals of 0.1. Error growth occurs sporadically, with three periods of rapid growth (commencing at $t = 0.1$, 0.9 and 1.6) being separated by short intervals of decay. The lower graph indicates the correlation of the two bred vectors with the leading orthogonal vector, and the vertical jumps in the lower graph indicate the times when the orthogonal vectors are reordered (at $t = 0.6$, 0.9 , 1.4 and 2.2). During the first and third periods of rapid error growth, the bred vectors converge rapidly to the leading orthogonal vector, but in the central period from $0.7 \leq t \leq 1.2$, they do not perform so well and both decay sharply even though the

orthogonal vectors show that there is a direction in which errors will be growing.

[Figure 1 here. Caption: A comparison of bred vectors (BV1, BV2, dotted and dashed lines respectively) and orthogonal bred vectors (OV1, OV2, thick and thin solid lines respectively). The upper panel (a) shows their lengths, and the lower panel (b) the correlation between each of the BV and the leading orthogonal vector.]

One obvious question is whether the high growth factor generated by the orthogonal pair is actually a realistic estimate of error growth in the model. After all, the second orthogonal vector has a mean growth factor of less than unity and therefore might be expected to shrink to a negligible size in the probability distribution function (pdf) of the model state. This is tested by use of a 4DVAR system (using an adjoint model) for assimilating observations into the model in an identical twin test. In these simple experiments, observations were taken for all 3 model variables at an interval of 0.1 time units, with independent Gaussian observation errors of size 2. The analysis interval was also 0.1 time units, using the end of the previous analysis interval as the start point for the next assimilation. Under these conditions, the RMS error at the endpoint of the analysis was 1.47, and the forecast error at 0.1 time units was 2.17 (all model results are presented in terms of the L_2 norm). The ratio of 1.48 is in good agreement with the growth factor generated by the orthogonal bred vectors, and significantly exceeds that of the standard bred vectors.

A further test can be performed by comparing the ability of a paired perturbation to enhance forecast skill. It is well known that the mean of a pair of forecasts with opposite perturbations will often outperform the control (unperturbed forecast) due to the nonlinear error saturation (eg Toth & Kalnay, 1997, section 2). This is found to hold true both for the bred vectors and orthogonal bred vectors: the improvements in this setup are extremely small (as we are far from saturation) but nevertheless consistent over repetition of the Monte Carlo tests. A standard bred vector pair improves the forecast skill by around 1.7×10^{-3} , the fastest-growing of 3 improves that to 2.3×10^{-3} (and the fastest of 20 manages 3.2×10^{-3}) but the leading vector from an orthogonal pair comfortably outperforms all of these with an improvement of 4.2×10^{-3} . Various methods of combining the bred vectors were tried but none was found to perform as well as simply using the fastest growing. The same was true for the orthogonal vectors. However this does not rule out the possibility that more sophisticated approaches might fare better.

Perhaps the most compelling evidence of all is simply to compare the direction of the true forecast error with that of the (leading) bred vector. A typical bred vector has a mean absolute correlation of 0.74 with the true forecast error (implying an angle of more than 42° between the two vectors), this rises to about 0.76 (41°) for the largest of 3 independent bred vectors, but the leading vector of an orthogonal pair achieves a correlation in excess of 0.80 (an angle of less than 37°).

These tests were repeated under different breeding conditions, with an observa-

tional interval of 0.3 time units and observational errors of size 3. In this case, the analysis error is 2.99 in the L_2 norm, and the forecast error is 8.93, representing an error growth factor of 2.99 over this time interval. Standard bred vectors of the same size have a growth factor of 2.42 over the same time interval, the largest of 3 bred vectors averages 2.47 and an orthonormal pair achieves a growth factor of 3.12. The forecast skill of paired perturbations is enhanced by scaling them back somewhat from the size of the analysis error (a factor of 0.55 was found by trial and error to give optimal results for both the standard and orthogonal breeding methods, the reason for this is not clear) and again the orthogonal method outperformed the standard breeding, reducing the forecast error by a substantial 0.80 (almost 9% of the forecast error) versus the consistently (albeit marginally) poorer figure of 0.77 from the best of 3 bred vectors.

The slight overestimate of error growth provided here by the orthogonal method can be attributed to the use of a fixed perturbation scale in each orthogonal direction. In reality, the expected shape of the analysis error will be non-spherical, and the dimensions of the quasi-ellipsoidal pdf will depend on the recent historical growth along its axes. For example, with isotropic observation errors of magnitude 2, and a forecast error of magnitude 1.4 in a particular direction, the analysis error in this direction should have magnitude 1.15. In an orthogonal direction where the forecast error is only 1, the analysis error will be 0.89. Rescaling each orthogonal perturbation to a fixed size (matching the overall expected error) therefore

artificially inflates the uncertainty along the second (and subsequent) orthogonal directions, and so will tend to exaggerate the likelihood of their magnitude exceeding that of the current dominant vector in the subsequent breeding cycle. It may be worthwhile to scale each orthogonal perturbation according to its expected contribution to analysis error, but this is not investigated here. In contrast, the standard bred vector method will tend to generate vectors which (when compared to the real analysis errors) have too small a component in the second and subsequent directions, since these are all regularly rescaled down according to the larger growth factor of the dominant direction of error growth. 'Stochastic' small-scale forcing due to local convection will prevent these dimensions from completely collapsing in NWP models, but it has also been found that adding artificially generated random perturbations onto the bred vectors can significantly improve their performance (Corazza *et al.*, 2002), by seeding the growth of perturbations in directions which are perpendicular to the orientation of the unperturbed bred vectors. The Ensemble Transform Kalman Filter method of Bishop *et al.* (2001) aims to generate perturbations that are a more accurate representation of analysis errors, by treating the ensemble members as a Monte Carlo sample of the error covariance and transforming them via the Kalman filter equations. This method also generates perturbations with a greater spread than both locally and globally rescaled bred vectors and finds faster maximal growth rates. However Wang & Bishop (2003) found that it was vulnerable to spurious sample correlations when

used with a small ensemble size, with 16 vectors clearly outperforming 8 in their experiments with a global atmospheric model. The orthogonalisation procedure described here is more akin to directly calculating the leading (local, finite size) Lyapunov vectors and so may be more useful for very small ensembles.

In summary, it appears that on all counts, the orthogonal method is slightly but clearly superior to the independently bred vectors for the Lorenz63 model. The orthogonalised vectors are less highly correlated, and therefore do a better job of spanning the space of growing errors. They are more efficient at generating rapidly growing perturbations, and these provide a realistic estimate of forecast errors in a 4DVAR system.

3 Local versus global analysis

The above results were generated by a point model, where the concept of local analysis does not have any meaning. Furthermore, the independently bred vectors are highly correlated (their mean absolute correlation is around 0.8, and were it not for the random perturbations, they would collapse to a single value), so it is not perhaps so surprising that orthogonalisation can help us to model the error space more efficiently. In contrast, bred vectors for NWP models are quasi-orthogonal, with one set of experiments finding a mean absolute correlation of 0.27 implying around 13 degrees of freedom (Toth & Kalnay, 1997). So further

global orthogonalisation seems unlikely to give much, if any, benefit here.

However, much attention has focussed recently on a more local analysis of atmospheric models. There are both theoretical and practical reasons why this is attractive. In schemes such as the ensemble Kalman filter (Evensen, 1994), long-range correlations are generally small and therefore cannot be estimated accurately with a realistic ensemble size (van Leeuwen, 1999). In practice, filtering long-range covariances improves the performance of ensemble Kalman filter schemes, while also reducing the required ensemble size for good performance (eg Houtekamer & Mitchell, 1998). Furthermore, the local dimensionality of atmospheric dynamics is relatively low (Patil *et al.*, 2001), which can be exploited by efficient analysis schemes (Ott *et al.*, 2003). Therefore, it seems reasonable to consider whether the handful of locally orthogonal growing modes can be captured by a handful of locally orthogonalised bred vectors. For example, if one considers a simple system with 6 independent regions each containing 2 equally probable orthogonal growing modes, then there are 12 globally orthogonal growing Lyapunov vectors, but a pair of locally orthogonal bred vectors would automatically include all of the growing local perturbations. In contrast, if independent bred vectors are used, then 5 of them would only have a 67% chance of containing all of the local modes. A first attempt to consider the issues of local analysis was undertaken using two copies of the Lorenz63 model, as coupled together by J. A. Hansen (2003, personal communication). The resulting 6-variable model was originally designed to study

the behaviour of coupled systems with different time and space scales, with appropriate scaling factors applied to one of the two subsystems (eg Peña & Kalnay, 2003). However, we are here merely trying to simulate a system with spatially distinct but otherwise similar regions of error growth, so no scalings are applied and the two submodels are identical in every respect other than in their initial conditions.

Breeding experiments with this coupled model were performed using the same parameters as above. With weak or no coupling, the behaviour of the coupled system under standard breeding is somewhat interesting. The bred vectors do not consist of similar-sized perturbations over both submodels, but instead a typical perturbation is overwhelming contained within one half of the coupled model, with a negligible disturbance in the other half. The reason for this is that the growth of perturbations varies independently in each submodel. During an interval where one region has a faster average growth than the other, the perturbation for the slower-growing region is repeatedly scaled down faster than it grows. The result is that the perturbation does not look like a realistic analysis error, and its growth rate is approximately equal to that of the fastest of the two regions rather than any sort of global average. Toth & Kalnay (1997) address this type of problem, and indeed go rather further than this, by their use of regionally (smoothly) varying rescaling factors which ensure that the bred perturbation is of similar size to the typical analysis error on a regional and not merely global average scale. The fact

that the rescaling factor is different in some distant regions will not in itself affect the local performance (under the assumption that the short-term dynamics are local), however a spatially varying factor could introduce a small loss of balance and appears to marginally degrade the growth of vectors (Wang & Bishop, 2003). Of course this penalty is offset by the benefits of having a vector which is locally of similar magnitude to the analysis error.

When the rescaling and orthogonalisation is performed on a ‘regional’ basis (in this case, to each submodel independently), then, at least in the case of rather weak coupling, the performance of the various breeding techniques can be deduced directly from their application to the basic 3-variable model. So these experiments are not discussed further here.

Of course, this simple coupled model does not address the issues of spatial continuity that will be crucial in any realistic spatially resolved model. To investigate this further, in the following section we apply these ideas to a model with somewhat more meaningful spatial variation, the 40-variable model of Lorenz (1996).

4 Local orthogonality for the Lorenz96 model

The Lorenz96 model (Lorenz, 1996) is a 40-variable model which has been used by various authors as a low-order proxy for atmospheric prediction and assimilation studies (eg Lorenz & Emmanuel, 1998; Anderson, 2001). The model variables X_i

evolve according to

$$dX_i/dt = (X_{i+1} - X_{i-2})X_{i-1} - X_i + F$$

where the indices $1 \leq i \leq 40$ are arranged cyclically and F is a fixed forcing. The model is integrated via the 4th-order Runge-Kutta method with a time step of $\Delta t = 0.05$, and when $F = 8.0$ the model displays sensitive dependence on initial conditions. The rescaling interval was chosen as 0.2 time units (representing 1 ‘day’) with the local perturbation size chosen to average 0.5 units at each variable (around 10% of the climatological range).

Local rescaling and orthogonalisation is performed using a simple windowing method inspired by the work of Kalnay & Toth (1994), Patil *et al.* (2001) and Ott *et al.* (2003). For any bred vector $V_j = v_{i,j}$, $1 \leq i \leq 40$, the *local vector* at the point i is given by the $(2l + 1)$ points $\mathbf{v}_{i,j} = v_{m,j}$, $i - l \leq m \leq i + l$ (with cyclic indices as in the model, and l a fixed integer). At each rescaling time, the i th variable in each bred vector is processed according to the local vectors $\mathbf{v}_{i,n}$, $1 \leq n \leq k$ where k is the number of bred vectors. These local vectors are rescaled and orthogonalised according to the Gram-Schmidt procedure (ordered by the size of the local vectors), and then the central elements (only) are used as the i th variable of the new global model state vectors. So each position in the updated model state is generated according to a Gram-Schmidt orthogonalisation applied to the local vectors centred on that point, and the local vectors are always been filled with

the unprocessed model variables. A spatially varying rescaled size could also be implemented, using the same method as Toth & Kalnay (1997), but this has not been investigated here. In the limit as the window size increases to the full length of the vector, the method converges to global Gram-Schmidt orthogonalisation. The local nature of the implementation results in some spatial variability in the size of the local vectors, the smoothness of which depends on the window size, rescaling interval and local variation in growth. Furthermore, there is the possibility of a discontinuity in the rescaled vectors at the transition where their local growth rates cross over, and the ordering of the Gram-Schmidt orthogonalisation switches. However, if rescaling is frequent enough, the vectors will always be near to local orthogonality and therefore this projection term should hopefully remain small.

It is not entirely trivial to assess local growth rates in an objective and fair method. Since the locally rescaled perturbation size is not identical at all points (it cannot be unless the perturbation is periodic), some of the apparent growth may be due to the advection of a larger initial perturbation into a region where the initial perturbation was rather small. This effect adds some noise to the estimation of local growth rates in some of the experiments. A more sophisticated smoothing could perhaps reduce this problem. Some simple experiments with globally rescaled bred vectors (with and without orthogonalisation) provide a benchmark which does not suffer from this effect. Under the conditions described above, a typical

globally rescaled bred vector grows by a factor of 1.43. When an ensemble of independent bred vectors is used, the fastest growth factor increases to 1.49 (with 2 independent vectors), 1.52 (3 vectors), 1.59 (10 vectors). The mean absolute correlation between any two bred vectors is 0.82. When the bred vectors are globally orthogonalised, their performance is improved. The fastest growing rate is 1.59 with 2 orthogonal vectors, 1.67 (with 3 vectors), 1.87 (10 vectors). These results show a clear enhancement of maximal growth rate when compared to normal bred vectors. As importantly, the typical correlation is only 0.17 at the end of the integration time, indicating that these vectors are spanning a much broader space of growing perturbations.

For the local rescaling and orthogonalisation, a window size of 7 ($l = 3$) gave reasonable results and is used throughout ($l = 2$ and $l = 4$ gave similar results which are not presented here). The local variation in the rescaling causes a modest reduction in average growth rates, which Wang & Bishop (2003) also found when comparing global and locally rescaled bred vectors. The mean maximal local daily growth rates found in ensembles of various numbers of independent bred vectors are: 1.27 (1 vector), 1.41 (max of 2), 1.48 (3) 1.57 (5), 1.66 (10), and 1.75 (20 vectors). When the vectors are locally orthogonalised, the mean maximal growth rates are: 1.27 (1 vector), 1.55 (max of 2), 1.68 (3), 1.85 (5), and 1.95 (7). The typical local correlation is around 0.67 for the simple bred vectors, but only 0.3 in the locally orthogonalised ensemble rising to 0.4 at the end of the

integration interval. This separation of vectors may be a more significant gain than the increased growth found. The maximal local growth rates do not appear to converge as quickly as in the previous experiments. This is probably due to the advective effect described earlier, as these apparent local growth rates exceed those found by global breeding. Nevertheless, it is clear that the locally orthogonal vectors are generating more rapid growth and also spanning a wider space of local perturbations.

For the ensembles of 3 and more orthogonalised vectors, the mean minimum local growth rates are below unity. These decaying perturbations are not relevant to forecast and analysis errors. Even slowly growing modes will not be of great interest, since they will at best represent small Lyapunov vectors and will in practice probably be rather noisy. It is not easy to quantify noise in this simple model since the optimal perturbations are not at all smooth. However, since the forecast errors are dominated by the fastest growing perturbations, it may be more useful in practical applications to resample the leading orthogonal perturbations by treating the ensemble as several small independent orthogonal sets rather than attempting to generate a large mutually orthogonal set that samples slow growing modes. The extent to which the lower order orthogonal vectors remain useful and noise-free will no doubt be case-dependent, but for example, within the context of this experiment, two parallel ensembles of 3 orthogonal vectors generated mean maximal growth rates equal to that of 7 mutually orthogonal vectors.

An interesting and perhaps important point to note is that the performance of the orthogonalised vectors is hardly altered if the local perturbations are orthogonalised in a fixed canonical order rather than according to size. This is true even in the 3 equation Lorenz63 model. In this case, the first two orthogonal bred vectors have near-identical growth factors of 1.18, the same as a simple bred vector (of course the first orthogonal vector here now is precisely equivalent to a simple bred vector). However the orthogonalisation means that these two vectors span a wide space of growing perturbations, with the (instantaneously) fastest growing averaging a growth of 1.43 and the slower averaging 0.93, in exact agreement with the results when the orthogonalisation is ordered according to size. If a third orthogonal vector is added, it has a growth factor of 0.43 and is so well differentiated from the first two that it does not affect their results. These results carry over to the local orthogonalisation of the Lorenz96 model, which also gives essentially identical results for maximal growth rates whether the local vectors are ordered by size or given a fixed order. For example, when the vectors are treated in a fixed order, we find mean maximal growth factors of 1.52 (with 2 vectors) 1.67 (3 vectors), 1.87 (5), and 1.98 (7) with no ordering prior to orthogonalisation. The local correlation of the vectors is even slightly less than in the size-ordered case, with the mean correlation rising from 0.15 after processing to 0.35 at the end of the integration period. Since this variation of the method eliminates the possibility of discontinuities through changing the ordering of orthogonalisation, it may

be more appropriate to use it in applications with realistic models. In either case, the maximum local growth rates are significantly enhanced by orthogonalisation, and the spread of growing perturbations is greatly improved. Both of these factors have the potential to enhance the use of bred vectors in forecasting and analysis schemes.

In the absence of a high-quality analysis scheme, no attempt was made to investigate the extent to which the various perturbation growth factors actually match the growth factors of forecast errors. However, the results from the Lorenz63 model are highly suggestive that the higher growth factors are likely to be more appropriate and that the standard breeding method has a modest but systematic bias towards underestimating the leading growth factors. Indeed, since the purpose of the bred vector method is specifically to generate the fastest growing finite-sized perturbations, the ability of locally orthogonalised bred vectors to find faster growth rates (for similar sized perturbations) is in itself sufficient demonstration of this slight weakness in the breeding method. It remains to be seen, however, whether the improvement demonstrated here can be carried over into the more complex setting of NWP models.

5 Conclusions and future work

Breeding provides an efficient and effective way of generating rapidly growing perturbations, and independent bred vectors in NWP models are globally quasi-orthogonal without the need for further constraints. However, short-term error growth is a local phenomenon, and when examined locally, bred vectors are frequently found to occupy a low dimensional subspace. Even though the local space of growing errors is also of low dimension, independent bred vectors may not span this space efficiently, and moreover even when many vectors are used they tend to sample a lower dimensional subspace of the growing errors. In tests with the 3-variable Lorenz63 model, orthogonalising the bred vectors enables them to capture the fastest growing perturbations more effectively, which also results in a more accurate prediction of forecast skill and a greater benefit due to ensemble averaging. With the spatially resolved 40-variable Lorenz96 model, local orthogonalising again generated faster growing perturbations with much lower local correlation. It would seem worthwhile to investigate to what extent these results can be carried over into more complex models including realistic NWP systems. It seems unlikely that a very large ensemble could be usefully treated in this way. The vectors at the tail of a long chain of orthogonalisation will become increasingly noisy, and since they are unlikely to represent rapidly growing modes, can in any case be omitted without penalty. Perhaps even an intermediate size ensemble would be better treated as a few small independent sets of locally orthogonalised pairs or

triples rather than one larger mutually orthogonalised set. Further experiments with more realistic models will certainly be required.

Another potential use of these ideas is to create a data assimilation scheme using orthogonal bred vectors to estimate the covariance matrix. The ensemble Kalman filter (Evensen, 1994) is a well-known and increasingly widely used assimilation system, but it is at the margin of practical applicability due to the large number of model runs required to sample the covariance matrix adequately. There have been several methods developed which aim to reduce the calculation size to more acceptable levels. Some of these are based on reduced rank approximations (Cohn & Todling, 1996; Heemink *et al.*, 1997; Pham *et al.*, 1998) which require the use of linearised models. These methods can suffer difficulties in the case of more highly nonlinear problems, and also require a certain amount of development for the linear model. Taking a somewhat different approach, Kalnay & Toth (1994) and Corazza *et al.* (2002) have used bred vectors to enhance a statistically-based covariance matrix with great effect, but since, as described above, the bred vectors are not in principle drawn from exactly the same pdf as the model state error, there are limits to how far this approach can be expected to go. It seems reasonable to hope that (locally) orthogonal bred vectors could help to reduce or perhaps even overcome the problems encountered in both of these strategies.

6 Acknowledgements

The author is grateful for many helpful comments from Eugenia Kalnay and Jim Hansen, and also to Jim Hansen for the coupled version of the Lorenz model.

References

- Anderson, J. L., 2001: An ensemble adjustment Kalman filter for data assimilation, *Monthly Weather Review*, **129**, 2884–2902.
- Bishop, C. H., B. J. Etherton & S. J. Majumdar, 2001: Adaptive sampling with the ensemble transform Kalman filter. Part 1: Theoretical aspects, *Monthly Weather Review*, **129**, 420–436.
- Cohn, S. E. & R. Todling, 1996: Approximate data assimilation for stable and unstable dynamics, *Journal of the Meteorological Society of Japan*, **74**(1), 63–75.
- Corazza, M., E. Kalnay, D. J. Patil, E. Ott, J. Yorke, I. Szunyogh & M. Cai, 2002: Use of the breeding technique in the estimation of the background error covariance matrix for a quasigeostrophic model, in *AMS Symposium on Observations, Data Assimilation and Probabilistic Prediction*, pages 154–157, Orlando, Florida.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *Journal of Geophysical Research*, **99**(C5), 10143–10162.
- Heemink, A. W., K. Bolding & M. Verlaan, 1997: Storm surge forecasting using

- Kalman filtering, *Journal of the Meteorological Society of Japan*, **75**(1B), 305–318.
- Houtekamer, P. L. & H. L. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique, *Monthly Weather Review*, **126**, 796–811.
- Kalnay, E., M. Corazza & M. Cai, 2002: Are bred vectors the same as lyapunov vectors?, in *AMS Symposium on Observations, Data Assimilation and Probabilistic Prediction*, pages 173–177, Orlando, Florida.
- Kalnay, E. & Z. Toth, 1994: Removing growing errors in the analysis, in *Preprints, 10th AMS conference on Numerical Weather Prediction*, pages 212–215.
- Kincaid, D. R. & E. W. Cheney, 1990: *Numerical Analysis: The Mathematics of Scientific Computing*, Brooks/Cole Publ.
- Lorenz, E. N., 1963: Deterministic non-periodic flow, *Journal of the Atmospheric Sciences*, **20**, 130–141.
- Lorenz, E. N., 1996: Predictability: a problem partly solved, in *Proc. ECMWF seminar on predictability*, pages 1–18, ECMWF, Reading, UK.
- Lorenz, E. N. & K. A. Emmanuel, 1998: Optimal sites for supplementary weather observations: simulation with a small model, *Journal of the Atmospheric Sciences*, **55**, 399–414.

- Ott, E., B. R. Hunt, I. Szunyogh, M. Corazza, E. Kalnay, D. J. Patil, J. A. Yorke, A. V. Zimin & E. J. Kostelich, 2003: Exploiting local low dimensionality of the atmospheric dynamics for efficient ensemble Kalman filtering, *Submitted to MWR*.
- Patil, D. J., B. R. Hunt, E. Kalnay, J. A. Yorke & E. Ott, 2001: Local low dimensionality of atmospheric dynamics, *Phys. Rev. Lett.*, **86**, 5878–5881.
- Peña, M. & E. Kalnay, 2003: Separating fast and slow modes in coupled chaotic systems, *Geophysical Research Abstracts*, **5**, 07540.
- Pham, D. T., J. Verron & M. C. Roubaud, 1998: A singular evolutive extended Kalman filter for data assimilation in oceanography, *Journal of Marine Systems*, **16**, 323–340.
- Toth, Z. & E. Kalnay, 1993: Ensemble forecasting at NMC: the generation of perturbations, *Bulletin of the AMS*, **74**(12), 2317–2330.
- Toth, Z. & E. Kalnay, 1997: Ensemble forecasting at NCEP and the breeding method, *Monthly Weather Review*, **125**, 3297–3319.
- van Leeuwen, P. J., 1999: Comment on: “Data asimilation using an ensemble Kalman filter”, *Monthly Weather Review*, **127**, 1374–1377.
- Wang, X. & C. H. Bishop, 2003: A comparison of breeding and ensemble transform

kalman filter ensemble forecast schemes, *Journal of the Atmospheric Sciences*, **60**, 1140–1158.

Yamane, S. & S. Yoden, 2001: Finite-time evolution of small perturbations superimposed on a chaotic solution: experiments with an idealized barotropic model, *Journal of the Atmospheric Sciences*, **58**, 1066–1078.

Figure Captions

Figure 1: A comparison of bred vectors (BV1, BV2, dotted and dashed lines respectively) and orthogonal bred vectors (OV1, OV2, thick and thin solid lines respectively). The upper panel (a) shows their lengths, and the lower panel (b) the correlation between each of the BV and the leading orthogonal vector.

A Comparison of Bred and Orthogonal Bred Vectors

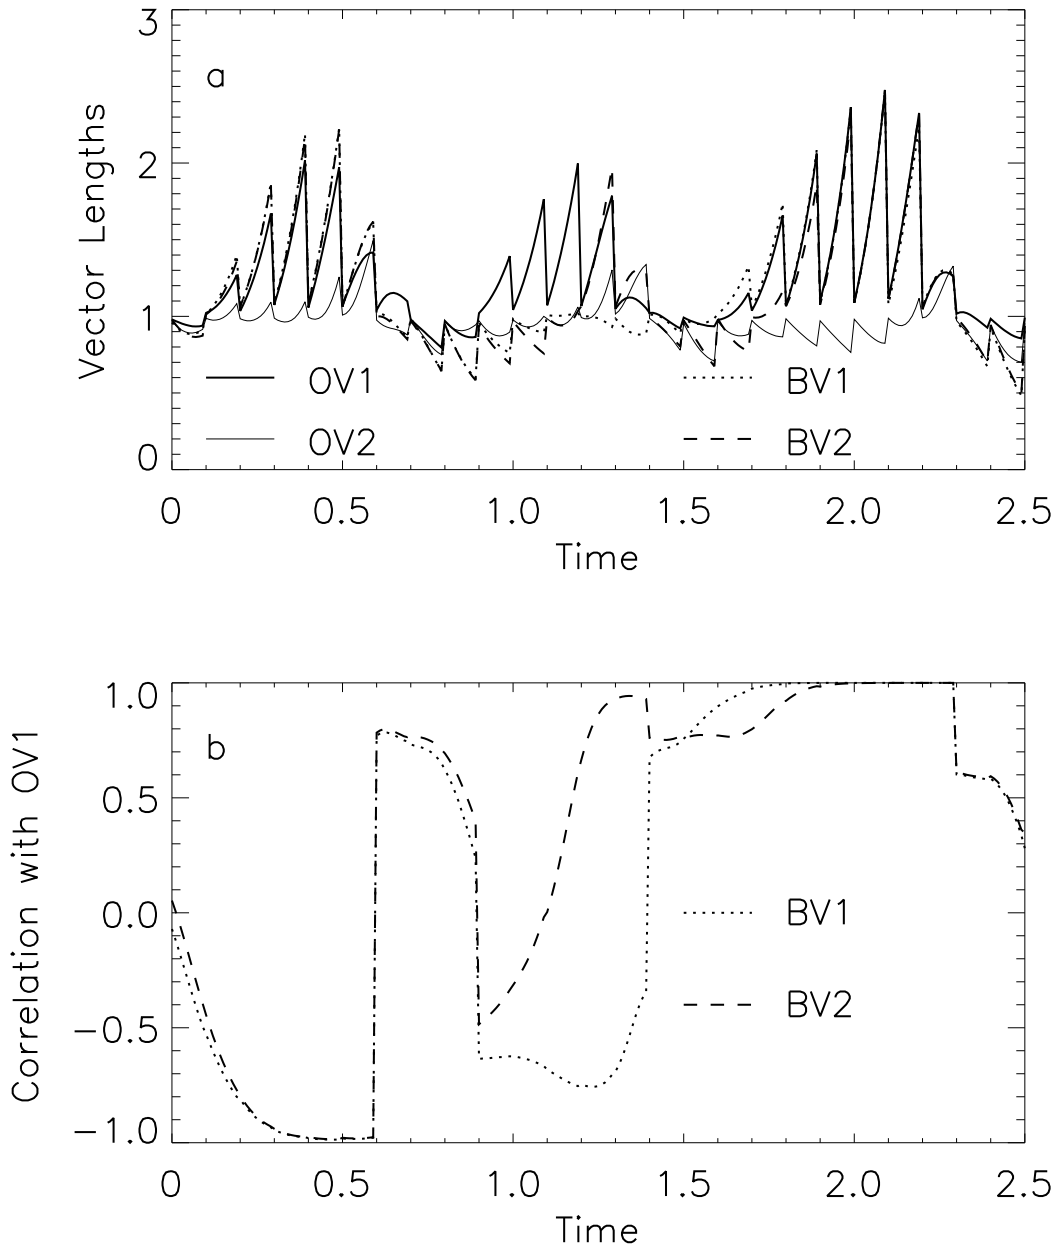


Figure 1: A comparison of bred vectors (BV1, BV2, dotted and dashed lines respectively) and orthogonal bred vectors (OV1, OV2, thick and thin solid lines respectively). The upper panel (a) shows their lengths, and the lower panel (b)

the correlation between each of the BV and the leading orthogonal vector.