Tellus (2000), 52A, 391–411

TELLUS
ISSN 0280–6495

Examination of targeting methods in a simplified setting

By C. A. REYNOLDS[1]*, R. GELARO[1] and T. N. PALMER[2], [1]*Naval Research Laboratory, 7 Grace Hopper Avenue, Monterey, CA 93943-5502, USA;* [2]*European Centre for Medium-Range Weather Forecasts, Shinfield Park, Reading, Berkshire, RG2 9AX, UK*

(Manuscript received 15 March 1999; in final form 10 January 2000)

ABSTRACT

The effectiveness of 2 methods for targeting observations is examined using a T21 L3 QG model in a perfect model context. Target gridpoints are chosen using the pseudo-inverse (the inverse composed of the first three singular vectors only) and the quasi-inverse or backward integration (running the tangent equations with a negative time-step). The effectiveness of a target is measured by setting the analysis error to zero in a region surrounding the target and noting the impact on the forecast error in the verification region. In a post-time setting, when the targets are based on forecast errors that are known exactly, both methods provide targets that are significantly better than targets chosen at random within a broad region upstream of the verification region. When uncertainty is added to the verifying analysis such that the forecast error is known inexactly, the pseudo-inverse targets still perform very well, while the backward integration targets are degraded. This degradation due to forecast uncertainty is especially significant when the targets are a function of height as well as horizontal position. When an ensemble-forecast difference is used in place of the inexact forecast error, the backward integration targets may be improved considerably. However, this significant improvement depends on the characteristics of the initial-time ensemble perturbation. Pseudo-inverse targets based on ensemble forecast differences are comparable to pseudo-inverse targets based on exact forecast errors. Targets based on the largest analysis error are also found to be considerably more effective than random targets. The collocation of the backward integration and pseudo-inverse targets appears to be a good indicator of target skill.

1. Introduction

It has been shown that even small initial condition errors can have a large impact on forecast errors (Rabier et al., 1996). Consequently, improvements in the initial analyses over relatively data sparse regions such as oceans may lead to significant improvements in forecasts over the continents downstream. Because of limited resources, much attention has been given to determining the location where analysis errors would have the largest impact on the forecast for a particular region or phenomenon, and targeting this location for additional observations. Targeted observation experiments have now been carried out over the Atlantic (FASTEX, Joly et al., 1999) and Pacific (NORPEX, Langland et al., 1999b).

Several different methods have been proposed for choosing the location of the extra observations. These include adjoint and singular vector (SV) methods (Palmer et al., 1998; Bergot et al., 1999; Langland et al., 1999a; Gelaro et al., 1999), the ensemble transform technique (Szunyogh et al., 1999), and the quasi-inverse or backward integration technique (Pu et al., 1999). Targeting studies have already demonstrated some successes. Langland et al. (1999a) have shown that observations within the target region have a larger impact on the forecast errors than observations outside

* Corresponding author.
e-mail: reynolds@nrlmry.navy.mil

the target region. Gelaro et al. (1999) have shown that most of the impact of the additional observations on the forecast error is due to the small part of the analysis increment that projects onto the leading SVs.

Both Pu et al. (1998) and Reynolds and Palmer (1998) point out that there are significant differences between the backward integration technique and adjoint or SV-based techniques. The adjoint and SV-based pseudo-inverse techniques (Buizza et al., 1997; Gelaro et al., 1998) are designed to find the fastest growing component of the analysis error. The backward integration technique, in contrast, is designed to find most of the full analysis error, not just the fastest growing part. (As currently applied in operational systems, the backward integration technique does not find the part of the error that decays due to dissipation.) The backward integration technique does not require the adjoint model, and is relatively inexpensive compared to the calculation of singular vectors, although it does require ensemble forecasts. Pu et al. (1998) compared the backward integration technique to the adjoint technique in the context of targeted observations in a post-time setting. They found that for an isolated initial analysis difference and strong error growth, both techniques successfully produced initial-time perturbations in the region of that initial analysis difference. The initial-time perturbations were horizontally collocated but had significantly different vertical structures. It is not clear that these initial-time perturbations would have been horizontally collocated if the initial analysis difference itself had not been confined to a relatively small region.

The pseudo-inverse and the backward integration techniques have been examined and compared in a related context using a simple T21L3 quasi-geostrophic model (Reynolds and Palmer, 1998; Reynolds, 1999). In these studies, the initial-time perturbations produced by these methods were used as analysis corrections, not as guidance for targeting observations. These studies point out contrasting characteristics of these two techniques. The pseudo-inverse method accurately finds the fastest growing component of the initial error when operating on the forecast error. This is usually only a small part of the full analysis error. Because this technique focuses on the fastest growing portion of the forecast error, the signal-to-noise ratio at final time is large, and thus the technique is relatively insensitive to uncertainty in the verifying analysis at final time. However, this technique requires the calculation of SVs, which can be expensive for larger models, and requires the choice of a metric for measuring error growth.

The backward integration was found to be a good approximation to the inverse of the forward tangent propagator in a simple QG model in Reynolds and Palmer (1998). This technique gives a very good estimate of the full analysis error when the forecast error is known exactly. However, when realistic uncertainty is added to the verifying analysis, the erroneous projection of this uncertainty onto the decaying SVs amplifies when integrated backwards in time, resulting in an unrealistically large estimate of the initial analysis error. This can lead to a "corrected" analysis that is actually further from the true state than the uncorrected analysis.

Given the different characteristics of these two methods, it is of interest to examine how well these techniques work as targeting tools. In this study, the two targeting methods are not only compared to each other, but their effectiveness is compared to randomly chosen targets, and targets based on some features of the mean flow. The impact of uncertainty in the verifying analysis is examined. Also, post-time techniques (based on the forecast error) are compared to real-time techniques (based on SVs or ensemble differences).

The study is conducted in a very idealized manner, and as such is useful for qualitative (as opposed to quantitative) comparisons with more complex systems. It is not designed to address issues such as the impact of model errors, target sampling, or the assimilation of targeted observations. Instead, the focus is on estimating the optimality of these techniques in providing targets given a perfect model and assuming the additional observations effectively set the analysis error to zero in the target vicinity. The use of the simple model allows for the economical computation of many forecast integrations for each case, which in turn allows for a robust assessment of the optimality of the targets.

A brief discussion of the pertinent linear theory is given in Section 2. The model and method used are described in Section 3. The results are presented in Section 4, and a summary and conclusions are given in Section 5.

2. Background

For a perfect model and linear error growth

$$\boldsymbol{L}\boldsymbol{e}_{\mathrm{a}} = \boldsymbol{e}_{\mathrm{f}}, \tag{1}$$

where $\boldsymbol{L}$ is the forward tangent propagator (tangent linear matrix) of the model, $\boldsymbol{e}_{\mathrm{a}}$ is the error at initial time and $\boldsymbol{e}_{\mathrm{f}}$ is the forecast error. From (1),

$$\boldsymbol{e}_{\mathrm{a}} = \boldsymbol{L}^{-1}\boldsymbol{e}_{\mathrm{f}}, \tag{2}$$

which indicates that if the inverse of the forward tangent propagator is known, then operating on the forecast error with $\boldsymbol{L}^{-1}$ produces the analysis error. For large models, $\boldsymbol{L}^{-1}$ cannot be represented explicitly. Pu et al. (1997) have proposed running the tangent equations with a negative time step as an approximation to $\boldsymbol{L}^{-1}$. They also reverse the sign of the dissipation terms to maintain computational stability. They refer to this method as the quasi-inverse (referred to as the backward integration in the present study). Reynolds and Palmer (1998) showed that this is a good approximation to the true inverse in a quasi-geostrophic model even when negative dissipation terms are used (as long as the magnitude of the dissipation is not large). Wang et al. (1997) have employed this technique in the context of data assimilation. They suggest that the dissipation terms in a complex atmospheric model are sufficiently strong such that $\boldsymbol{L}$ becomes effectively singular. Thus, reversing the sign of the dissipation terms is necessary when approximating $\boldsymbol{L}^{-1}$. This technique has been used for targeted observation purposes (Pu et al., 1998, 1999) by first applying a mask to the forecast error, which effectively sets the forecast error to zero outside the verification region, and then operating on this field with the backward integration.

Alternatively, Buizza et al. (1997) have proposed using the pseudo-inverse to diagnose the fastest growing components of the analysis error. The pseudo-inverse requires the calculation of the SVs of $\boldsymbol{L}$. The SVs of $\boldsymbol{L}$ are the eigenvectors of $\boldsymbol{L}^*\boldsymbol{L}$, where $\boldsymbol{L}^*$ is the adjoint of $\boldsymbol{L}$ with respect to a specific metric. The SVs are actually calculated for $\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2}$ where $\boldsymbol{K}$ is a matrix such that for a state vector $\boldsymbol{x}$, $\boldsymbol{x}^{\mathrm{t}}\boldsymbol{K}\boldsymbol{x}$ gives the measure of perturbation growth, or metric. $\boldsymbol{L}$ is expressed in terms of SVs and singular values (see Molteni and Palmer, 1993) as such;

$$\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2} = \boldsymbol{V}_{\boldsymbol{K}}\Sigma_{\boldsymbol{K}}\boldsymbol{U}_{\boldsymbol{K}}^{\mathrm{T}}, \tag{3}$$

where $\boldsymbol{U}_{\boldsymbol{K}}(\boldsymbol{V}_{\boldsymbol{K}})$ is a matrix with columns that are the initial-time (final-time) SVs of $\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2}$, and $\Sigma_{\boldsymbol{K}}$ is a diagonal matrix whose elements are the singular values of $\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2}$, ranked in order of descending magnitude. The first column of $\boldsymbol{U}_{\boldsymbol{K}}$ represents the fastest growing perturbation to the specific trajectory considered, which evolves into the first column of $\boldsymbol{V}_{\boldsymbol{K}}$ and has an amplification factor given by the first singular value.

The inverse of the forward tangent propagator may be expressed in terms of the same singular values and SVs as above;

$$\boldsymbol{K}^{1/2}\boldsymbol{L}^{-1}\boldsymbol{K}^{-1/2} = \boldsymbol{U}_{\boldsymbol{K}}\Sigma_{\boldsymbol{K}}^{-1}\boldsymbol{V}_{\boldsymbol{K}}^{\mathrm{T}}. \tag{4}$$

The pseudo-inverse method uses $\boldsymbol{L}^{-1}$ as described in (4), but composed of only a limited subset of the fastest growing SVs in order to find the fastest growing components of $\boldsymbol{e}_{\mathrm{a}}$. The pseudo-inverse $\boldsymbol{L}_{\boldsymbol{K},\boldsymbol{n}}^{-1}$ is defined as

$$\boldsymbol{L}_{\boldsymbol{K},\boldsymbol{n}}^{-1} = \boldsymbol{U}_{\boldsymbol{K},\boldsymbol{n}}\Sigma_{\boldsymbol{K},\boldsymbol{n}}^{-1}\boldsymbol{V}_{\boldsymbol{K},\boldsymbol{n}}^{\mathrm{T}}, \tag{5}$$

where $\boldsymbol{U}_{\boldsymbol{K},\boldsymbol{n}}(\boldsymbol{V}_{\boldsymbol{K},\boldsymbol{n}})$ is a rectangular matrix with columns composed of the first n initial-time (final-time) SVs of $\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2}$ and $\Sigma_{\boldsymbol{K},\boldsymbol{n}}$ is a diagonal matrix composed of the first n singular values of $\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2}$. The fastest growing part of the initial error can then be defined as

$$\boldsymbol{e}_{\mathrm{a\,fast}} = \boldsymbol{K}^{-1/2}\boldsymbol{L}_{\boldsymbol{K},\boldsymbol{n}}^{-1}\boldsymbol{K}^{1/2}\boldsymbol{e}_{\mathrm{f}}, \tag{6}$$

for small n. For larger models, the first n SVs can be found through an iterative technique based on the Lanczos or Arnoldi algorithms applied to the tangent equations of the model. Note that the pseudo-inverse $\boldsymbol{L}_{\boldsymbol{K},\boldsymbol{n}}^{-1}$ is a function of the metric $\boldsymbol{K}$. However, when all the singular vectors are included, $\boldsymbol{L}_{\boldsymbol{K},\boldsymbol{n}}^{-1}$ becomes $\boldsymbol{K}^{1/2}\boldsymbol{L}^{-1}\boldsymbol{K}^{-1/2}$, and (6) becomes equivalent to (2) (i.e., the pseudo-inverse is a function of the metric while the full inverse is not).

In this study, a kinetic energy (KE) metric is used. In this case $\boldsymbol{K}$ is a diagonal matrix defined such that $\boldsymbol{x}^{\mathrm{t}}\boldsymbol{K}\boldsymbol{x}$ gives the kinetic energy of the system, and the SVs of $\boldsymbol{K}^{1/2}\boldsymbol{L}\boldsymbol{K}^{-1/2}$ are referred to as the KE SVs of $\boldsymbol{L}$. The appropriateness of an energy-based metric for these applications is discussed in Palmer et al. (1998). For targeted observation applications, a local projection operator can be applied to the calculation of the SVs (Palmer et al., 1998), following Barkmeijer (1992) and Buizza (1994). This is done by finding the KE singular values not of $\boldsymbol{L}$, but of $\boldsymbol{P}\boldsymbol{L}$, where the local projection operator $\boldsymbol{P} = \boldsymbol{S}^{-1}\boldsymbol{G}\boldsymbol{S}$. $\boldsymbol{S}$ is the matrix representation of the spectral to grid-point transform, and $\boldsymbol{G}$ is a diagonal matrix which sets to zero perturbations outside the region of interest. Applying this projection operator constrains the final-time SVs to optimize perturbation growth within the region of interest. However, there is no formal constraint on the locations of the SVs at initial time (Gelaro et al., 1999).

3. Approach

3.1. Model

The model used in this study is the quasi-geostrophic potential vorticity model described in Marshall and Molteni (1993). The model is run with a T21 truncation and has three levels corresponding to 800, 500 and 200 hPa. The model forcing is composed of specified source terms of PV that are spatially varying but temporally constant and correspond to a northern winter climatology. The model has three types of dissipative forcing: Newtonian relaxation of temperature between levels with a relaxation coefficient of 25 d^{-1}, an 800-hPa linear drag term that varies with topography from 3 to 1.5 d^{-1}, and horizontal scale-selective (∇^8) dissipation such that spherical harmonics of the time-varying potential vorticity with total wave number 21 are damped on a 2-day time scale. The model, with 1449 degrees of freedom, is well suited for this type of study, since it is complex enough to capture baroclinic synoptic scale processes important in forecast error growth, but small enough such that the SVs can be calculated explicitly and a large number of nonlinear integrations is feasible.

3.2. Experimental design

A schematic diagram of the experimental design is shown in Fig. 1. The 48-h integration of the nonlinear model representing the true atmospheric state starts at t_0 and ends at t_{48}. A second nonlinear integration starting from initial conditions slightly different from the first integration is taken as the forecast, with starting and ending states denoted by a_0 and f_{48}. The prescribed initial difference between a_0 and t_0 represents the analysis error, e_a. The exact forecast error, $e_{\text{f exact}}$, is given as $f_{48} - t_{48}$. Another perturbation, δ, represents

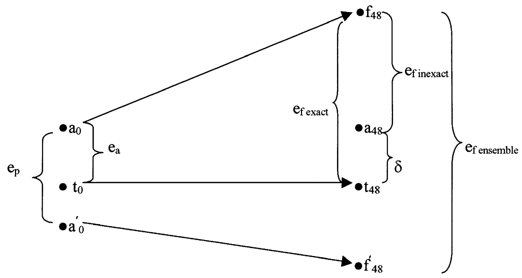

Fig. 1. Schematic representation of forecast experiments. The model state vectors are represented by dots and the 48-h nonlinear integrations are represented by arrows. The nonlinear integration representing the true atmosphere starts at t_0 and ends at t_{48}. The nonlinear integration representing the control forecast starts at a_0 and ends at f_{48}. The nonlinear integration representing an ensemble member forecast starts at a_0' and ends at f_{48}'. The initial-time analysis error (e_a), error in the verifying analysis (δ), and initial-time ensemble perturbation (e_p) are prescribed. The sizes of the brackets are not necessarily proportional to the magnitudes of the difference fields.

the error in the verifying analysis, a_{48}, such that $t_{48} - a_{48} = \delta$. Therefore, the observed, or inexact forecast error $e_{\text{f inexact}} = f_{48} - a_{48} = e_{\text{f exact}} + \delta$. The prescribed fields e_a and δ are taken from the differences between ECMWF OI and 3DVAR analyses. These fields have an average RMS geopotential height of 10 m, a broad KE spectral peak between total wave numbers 10 and 15, and more energy at 200-hPa and 500-hPa than at 800-hPa.

In addition to the above experiments, an ensemble member forecast is also simulated in a simplistic manner. A perturbation e_p is added to the initial-time analysis a_0 to give the ensemble initial state a_0', and a 3rd nonlinear integration is run from this state to produce the ensemble forecast f_{48}'. The difference between the control and ensemble forecast is given as $e_{\text{f ensemble}} = f_{48} - f_{48}'$. Three different types of perturbations are used for e_p, referred to as $e_{\text{p corr}}$, $e_{\text{p uncorr}}$, and $e_{\text{p lv}}$. The first type of initial-time ensemble difference, $e_{\text{p corr}}$ is an OI-3DVAR difference field. The OI-3DVAR difference fields used for $e_{\text{p corr}}$ and e_a are for different dates, but because of systematic differences between the OI and 3DVAR systems, the average global anomaly correlation between $e_{\text{p corr}}$ and e_a is 0.44. To create the second type of ensemble perturbation, $e_{\text{p lv}}$, an approximation to the global leading Lyapunov vector is used. The leading Lyapunov vector is approximated by the leading final-time singular vector optimized over the 40 days immediately preceding the analysis time (Reynolds and Errico, 1999). These vectors have a spectral peak in energy between total wave numbers 7 and 10, and the most energy at 200 hPa (see Fig. 7 from Reynolds and Errico, 1999). The ensemble perturbations generated at NCEP share some similarities with the leading Lyapunov vectors (Toth and Kalnay, 1997). Finally, ensemble perturbations are produced using randomly generated numbers but are constrained to have a prescribed KE spectra similar to that of the OI-3DVAR analysis differences used for e_{a}. Although constructed to have spectral characteristics similar to the initial-time analysis error e_{a}, $e_{\mathrm{p\,uncorr}}$ for any particular case does not, in general, resemble e_{a} (the average magnitude of the global anomaly correlation is 0.11). Please note that Fig. 1 is only a schematic, and the sizes of the brackets do not accurately reflect the relative magnitudes of the difference fields. The prescribed fields e_{a}, δ, and e_{p} are all of similar magnitude.

The goal of targeting methods is to provide the location where a reduction in analysis error (presumably through the addition of observations) will result in the largest forecast improvement over the verification region. In this study, the verification region is a box centered over Europe that extends from 37°N to 65°N and from 17°W to 34°E (Fig. 5).

The pseudo-inverse method is outlined here:

- Calculate 48-h KE SVs based on the nonlinear "forecast" trajectory using a local projection operator over the verification region.
- Form a pseudo-inverse of the tangent forward propagator, as in (5), using the first three SVs.
- Operate on the forecast error (or difference) field using the pseudo-inverse, as in (6), to find the component of the analysis error that maximizes error growth in the verification region.
- Use the initial-time perturbation found in step three to determine the target region.

The backward-integration method is outlined here:

- Approximate the full inverse of the tangent forward propagator by running the tangent equations with a negative time step (also reversing the sign of the dissipation terms).
- Operate on the forecast error (or difference) field (set to zero outside the verification region) using the backward integration, described in step one, to find the component of the analysis error that evolves into the forecast error in the verification region.
- Use the initial-time perturbation found in step two to determine the target region.

The initial-time perturbations are used to determine the location of the target region in two ways. The first way is to find the gridpoint location of the maximum of the vertically averaged squared potential vorticity of the perturbation. To mimic the effect of targeting in an idealized manner, the analysis error is set to zero within a radius of 10° latitude centered on the target grid-point at all three model levels. Between a radius of 10 and 20° from the target gridpoint, the analysis error is gradually increased from zero to the control value using a cosine function. These perturbations are referred to as the R10 analysis perturbations. The experiments are also repeated using smaller-scale analysis perturbations, such that the analysis error is set to zero within a radius of 3°, gradually increasing to the control value at a radius of 10°. These analysis perturbations are referred to as the R3 analysis perturbations (note this function is represented very coarsely at the 5.625 gridpoint resolution).

The second type of target is based on the largest value of the PV perturbation, and is a function of height as well as horizontal location. In this case, the analysis error is set to zero using a horizontal function as described above, but only at the level of the target gridpoint. The analysis error at the other two levels is left as is. The analysis perturbations based on the vertically averaged targets that set the analysis error to zero in the model column are referred to as the column analysis perturbations. The analysis perturbations based on the 1-level targets that set the analysis error to zero on one level only are referred to as the 1-level analysis perturbations.

These experiments are performed first using the exact forecast error, $e_{\mathrm{f\,exact}}$, referred to as the exact backward integration and exact pseudo-inverse experiments. They are then repeated using $e_{\mathrm{f\,inexact}}$, where uncertainty has been added to the verifying analysis. These experiments are referred to as inexact backward integration and inexact pseudo-inverse experiments. In an operational setting forecast errors are obviously not available, and ensemble differences are sometimes used to provide estimates of forecast error variance. Therefore similar experiments are conducted by operating on ensemble-forecast differences in place of the forecast errors, referred to as ensemble backward integration and ensemble pseudo-inverse experiments. All experiments are conducted for 20 different cases.

4. Results

4.1. *Case study*

For illustrative purposes, one case is examined in detail as an example of the procedure employed here. Figs. 2a, b show the 500-hPa geopotential height for t_0 and t_{48} respectively, denoted by the solid contours. Superimposed on these fields is the initial analysis error (e_a) and forecast error ($e_{f\,\mathrm{exact}}$), denoted by the shading. The forecast verification area is denoted by the box over Europe. As expected, the magnitude of the analysis error is significantly smaller than the forecast error. The main synoptic feature is a high-amplitude ridge over the Atlantic that retrogrades from the straight of Denmark to the southern tip of Greenland. The forecast error pattern over Europe indicates a phase error in the trough to the east of this ridge.

The vertically averaged squared PV of the exact backward integration initial-time perturbation is shown in Fig. 3. The vertically averaged squared PV of the exact pseudo-inverse initial-time perturbation is shown in Fig. 4. The locations of the maxima are marked on the figures. The concentric circles in Figs. 3, 4 give the scales of the R10 and R3 analysis perturbations, respectively. The analysis error is set to zero within the inner circle and increased to the control analysis error between the inner and outer circles as described in Section 3. When the forecast is rerun from the corrected analysis (R10 analysis perturbation), it is found that the target based on the exact backward integration reduces the forecast error over the European verification region by 37%. The target indicated by the exact pseudo-inverse methods reduces the forecast error by 41%.

These types of results are useful for comparing different targeting methods, but do not provide information on the optimality of individual targets. Toward this purpose, each gridpoint over the broad region extending from 85°W to 75°E and from 20°N to 65°N is considered as a possible target (i.e., the center of the analysis perturbation). A contour plot is produced showing the percent forecast error reduction obtained by setting the analysis error to zero in a column centered on each gridpoint. Forecast error is measured in terms of kinetic energy over the verification region, but results based on mean-square geopotential height are very similar. The contour plot based on the R10 analysis perturbations for this case is shown in Fig. 5. The European verification region is denoted by the small box. The results indicate that the best target would be in a region off the west coast of North Africa, producing forecast error reductions over Europe of greater than 45%. Also note there are some regions of negative contours, indicating that reducing the analysis error in these regions will actually *increase* the forecast error over the verification region. The contour plot based on the R3 analysis perturbations (not shown) has a structure similar to that shown in Fig. 5, but the maximum values are smaller.

Ideally, a targeting method will produce a target location that coincides with the maximum on this contour plot. The markers of the target gridpoints shown in Figs. 3, 4 are also shown in Fig. 5. The exact pseudo-inverse target and the exact backward integration target are quite close to this maximum.

In order to quantify the optimality of these targets, each gridpoint from 85°W to 23°E and from 20°N to 65°N (denoted by the large box in Fig. 5), is ranked from worst to best in terms of producing a forecast error reduction over the verification region. The size and location of this box are chosen such that it encompasses all the pseudo-inverse and backward integration targets obtained in the 20 cases. The rank is then normalized by the total number of gridpoints within this region. The worst target gridpoints have ranks near zero, and the best target gridpoint has a rank of one. Effective targeting methods should indicate gridpoints that are ranked among the top few percent of all gridpoints in terms of forecast error reduction. Viewed in this manner, the exact pseudo-inverse target has a rank of 0.96 and the exact backward integration target a rank of 0.91.

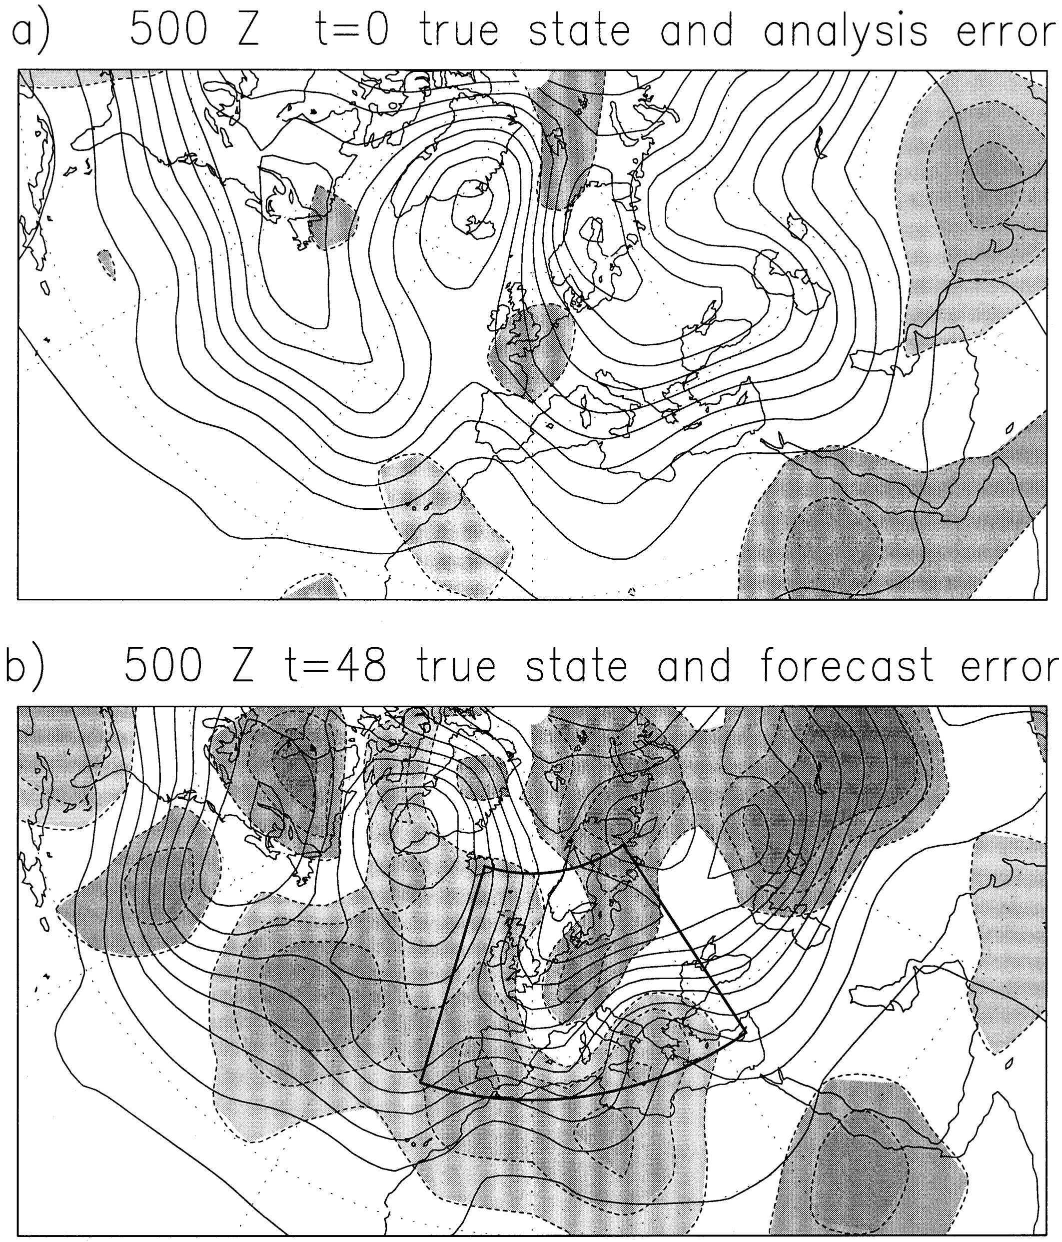

Fig. 2. The 500-hPa geopotential height fields for (a) t_0 and (b) t_{48} denoted by solid contours (contour interval of 60 m). The analysis error e_a (a) and forecast error $e_{\mathrm{f\,exact}}$ (b) are denoted by dashed contours [contour interval is 10 m, zero contour omitted, with values less than -10 m (greater than $+10$ m) indicated by light (dark) shading]. The box over Europe indicates the forecast verification region.

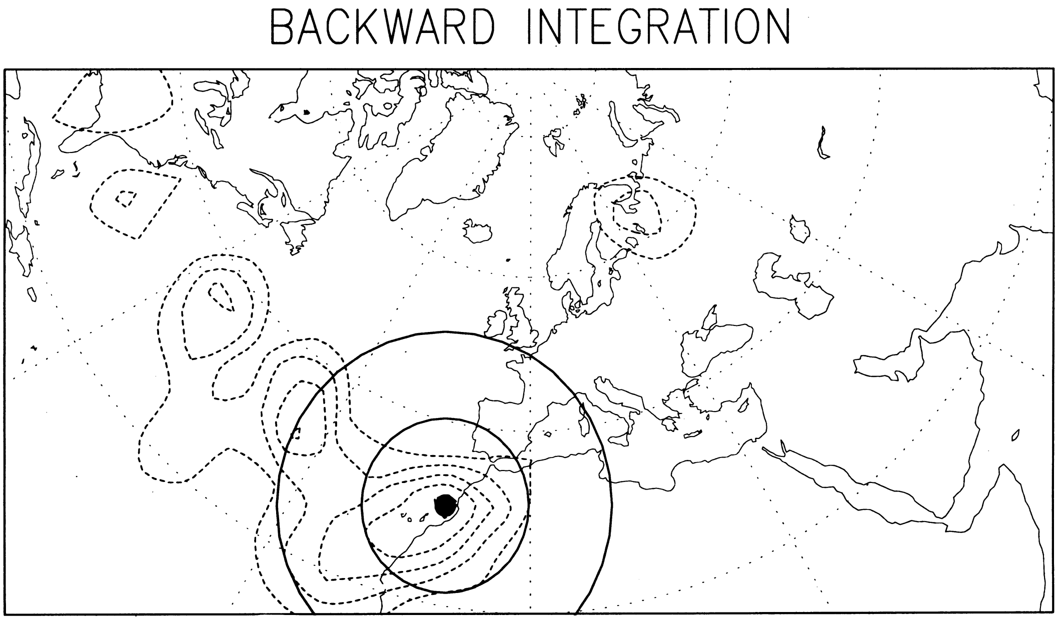

Fig. 3. The vertically averaged squared potential vorticity associated with the initial-time perturbation from the exact backward integration method (dashed contours). The location of the target gridpoint is denoted by the solid circle. The spatial extent of the R10 analysis perturbation is given by the concentric circles (solid contours, see text). Contour interval is arbitrary.

4.2. Composite results for post-time methods

Fig. 6 shows the ranks of the vertically averaged targets based on the exact and inexact backward integration (Fig. 6a) and the exact and inexact pseudo-inverse (Fig. 6b), using the R10 analysis perturbations. There is considerable case-to-case variability in the ranks of the targets. Also apparent is the significant degradation between the inexact and exact backward integration targets, which are collocated in only five of the 20 cases (and within one gridpoint of each other in another three cases). The average rank is 0.79 for the exact backward integration targets and 0.69 for the inexact backward integration targets. The pseudo-inverse targets are less sensitive to uncertainty in the verifying analysis, being collocated in 15 of the 20 cases (and within one gridpoint of each other in one case). The average ranks are 0.81 and 0.78 for the exact and inexact pseudo-inverse targets respectively. Note that a rank close to zero does not necessarily mean that the target is in an insensitive region of the atmosphere. On the contrary, it may mean that the target is in a sensitive region, but the analysis perturbation in that region is such as to actually increase the projection of the analysis error onto a fast-growing structure. This can happen, for example, when the initial analysis perturbation covers only part of a fast growing SV. Gelaro et al. (1999) and Langland et al. (1999a) report similar behavior in forecast experiments based on the FASTEX data set. Bergot et al. (1999) report on the importance of sampling a target well in order to obtain a significant beneficial impact.

Another way to examine the effectiveness of a particular targeting method is to produce histograms of the ranks of the targets. Fig. 7 contains histograms for the exact and inexact backward integration targets (Fig. 7a,b) and exact and inexact pseudo-inverse targets (Fig. 7c,d) for the R10 analysis perturbations. The ranks are divided into ten bins for each histogram. A near-perfect targeting technique would have targets occurring in the top bin (solid black, indicating a rank of 0.9 to 1.0), for all 20 cases, and no occurrences in the other bins. In contrast, target locations selected at random within the region 20°N–60°N and 85°W–23°E would produce a fairly flat histogram

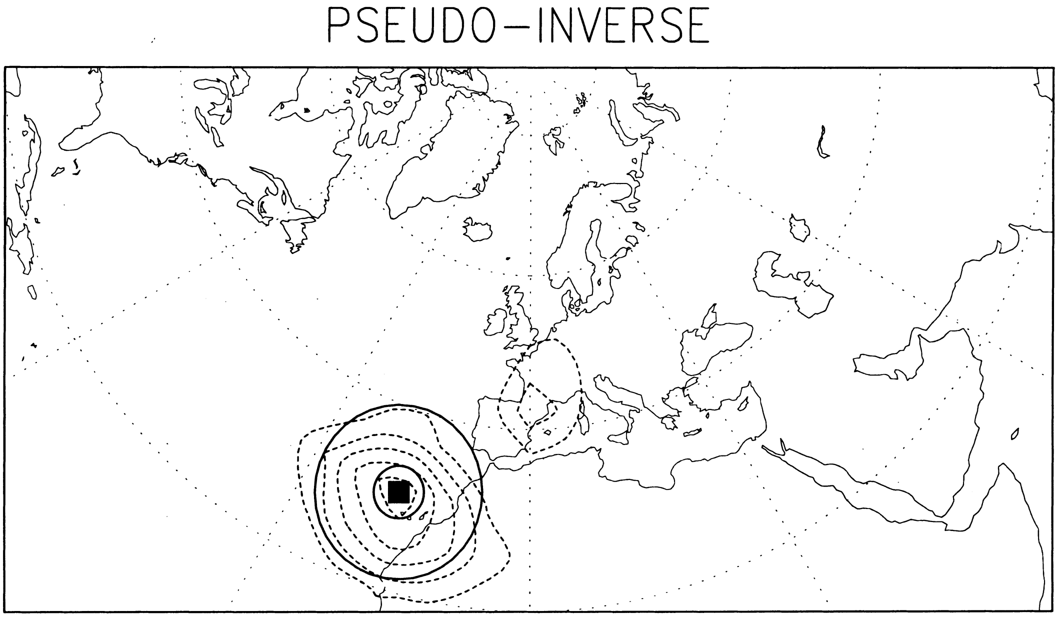

Fig. 4. The vertically averaged squared potential vorticity associated with the initial-time perturbation from the exact pseudo-inverse method (dashed contours). The location of the target gridpoint is denoted by the solid square. The spatial extent of the R3 analysis perturbation is given by the concentric circles (solid contours, see text). Contour interval is arbitrary.

with approximately two occurrences in each bin (for a sufficiently large sample, each bin would have almost the same number of occurrences). The dominance of the 0.9–1.0 bin in the four graphs indicates that all target methods are better than choosing targets at random. The backward integration and pseudo-inverse methods are comparable when based on exact forecast errors. The backward integration shows a significant degradation when uncertainty is introduced into the verifying analysis.

Fig. 8 is the same as Fig. 7 but for the R3 analysis perturbations. A comparison of Fig. 8 with Fig. 7 shows that the results for the rankings are similar for the two different analysis perturbation sizes (although the actual percent forecast error reductions, not shown, are larger for the R10 analysis perturbations). However, there are some differences. The backward integration targets have somewhat higher ranks with the R3 analysis perturbations. The pseudo-inverse targets have a few more occurrences at the very low ranks with the R3 analysis perturbations. Some sensitivity to analysis perturbation size is expected. If the scale of the analysis perturbation is small, then it is possible to inadequately sample the fast growing target structures, resulting in a suboptimal reduction in forecast error (or, in some cases, an increase in forecast error). On the other hand, if the scale of the analysis perturbation is large, its optimal position may not necessarily coincide with the extrema of the target structure. For example, if two fast growing SVs lie near each other, it may be best to center a small analysis perturbation on one of the SVs. On the other hand, if the analysis perturbation is sufficiently large, the best strategy may be to center the analysis perturbation between the two SVs. Further study is needed to determine an optimal sampling strategy that, in an operational setting, will be dependent upon the data assimilation scheme as well.

All results shown so far have been based on vertically averaged PV^2 targets and column analysis perturbations that set the analysis error to zero throughout the depth of the model. Targets that are a function of height have also been considered. In this case, the target is based on the horizontal and vertical location of the largest

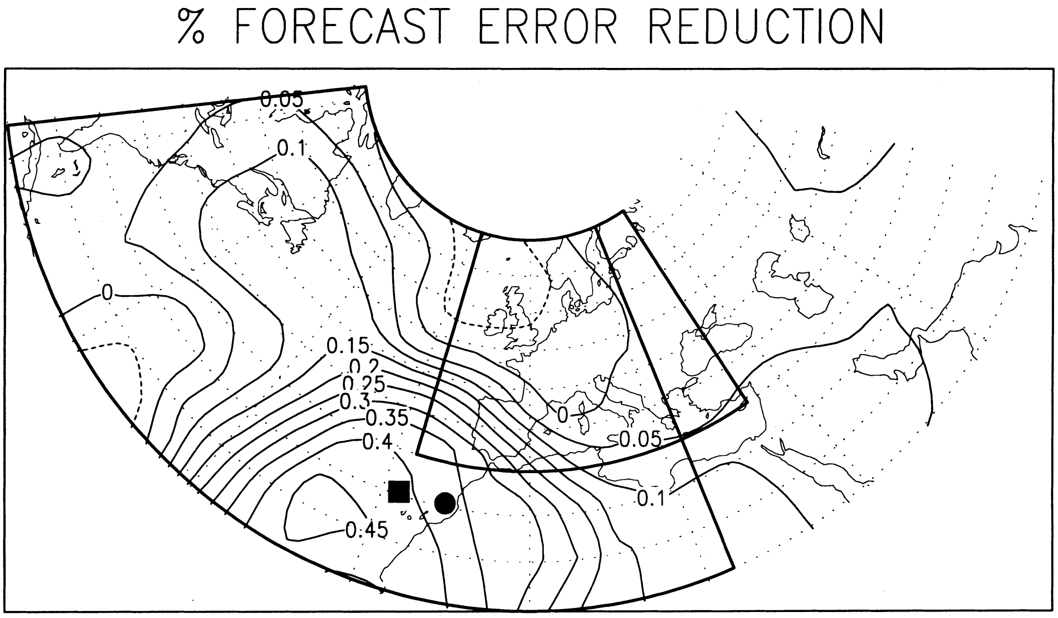

Fig. 5. Contour plot of the percent forecast error reduction obtained by centering the R10 analysis perturbation at each gridpoint. The small box over Europe indicates the verification region. The large box indicates the gridpoints that are included in the ranking processes. The target locations are indicated by the solid circle for the exact backward integration and the solid square for the exact pseudo-inverse.

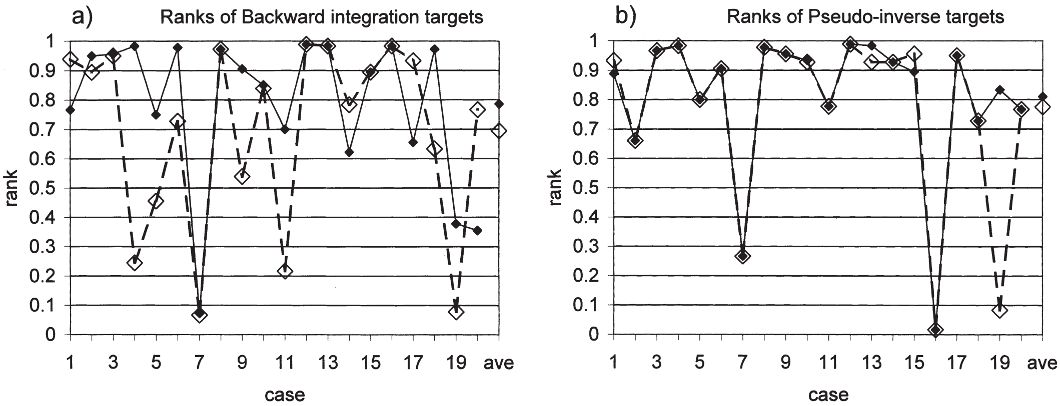

Fig. 6. The ranks for the vertically averaged targets based on the R10 analysis perturbations for all 20 cases for (a) the backward integration targets and (b) the pseudo-inverse targets. The thin solid line and solid diamond indicate the targets based on the exact forecast error. The thick dashed line and clear diamond indicate the targets based on the inexact forecast error. The symbols at the far right of the plot indicate the 20-case average values.

value of the initial-time PV perturbation. The analysis perturbation is confined to that vertical level such that the analysis error is set to zero only at that model level. Histograms of the ranks of these one-level targets for the R10 analysis perturbations are shown in Fig. 9. As with the vertically averaged targets shown in Fig. 7, the exact backward integration and exact pseudo-inverse targets are comparable (Fig. 9a,c). Both the exact and inexact pseudo-inverse targets

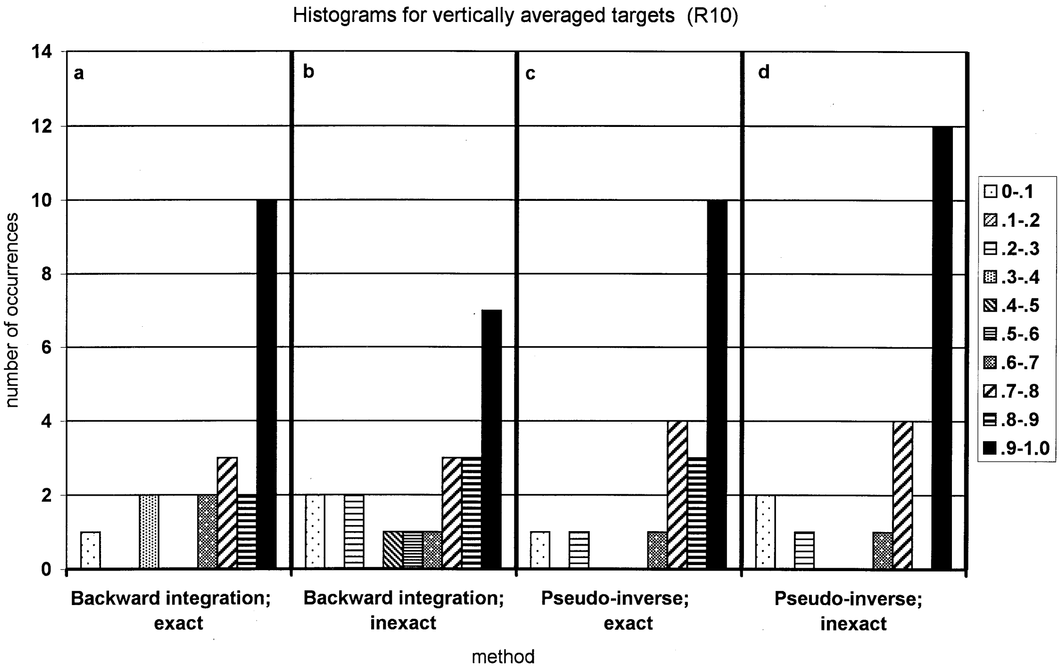

Fig. 7. Histograms of the rankings of the vertically averaged targets and R10 analysis perturbations for (a) the exact backward integration, (b) inexact backward integration, (c) exact pseudo-inverse and (d) inexact pseudo-inverse. The values of the ranks included in each bin are given by the key on the right.

perform very well, with only one case out of the 20 having a rank of less than 0.6. However, now the inexact backward integration targets are degraded to the point where they are not much more effective than targets chosen at random. The significant degradation between the exact and inexact backward integration targets also exists when using the R3 analysis perturbations (Fig. 10).

The average ranks of the targets, based on the R10 analysis perturbations, are stratified as a function of vertical level in Table 1. The inexact backward integration has many more targets at the 200-hPa level (14) than the other methods (2 to 4). These 200-hPa targets have a much lower average rank than those at 500 hPa for the backward integration cases. The average ranks of the exact and inexact backward integration 200-hPa targets are 0.4 and 0.42 respectively. The average ranks of the exact and inexact backward integration 500-hPa targets are 0.89 and 0.9 respectively. (The best target gridpoints are actually at 500 hPa in 15 of the 20 cases, and at 200 hPa and 800 hPa in three and two cases respectively.) It is interesting to note that the average ranks for the exact and inexact pseudo-inverse 200-hPa targets are 0.70 and 0.86, respectively. This indicates that the pseudo-inverse technique is not constrained to producing targets at the 500-mb level, and is able to provide effective targets at the 200-hPa level should they exist.

Fig. 11 shows the kinetic energy as a function of total wave number for the initial error, e_a, as well as the inexact backward integration and inexact pseudo-inverse initial-time perturbations, averaged for the 20 cases. (All fields have been normalized to have total kinetic energy of 1.0.) The pseudo-inverse perturbations have the most energy at the 500-hPa level, consistent with many previous studies concerning initial-time fast-growing SVs. Bergot et al. (1999) also find that that the most significant contributions for improving the forecast come from analysis improvements in the lower and middle troposphere. In contrast, the backward integration perturbations have the most energy at the 200-hPa level, consistent with the fact that the decaying SVs have a much larger

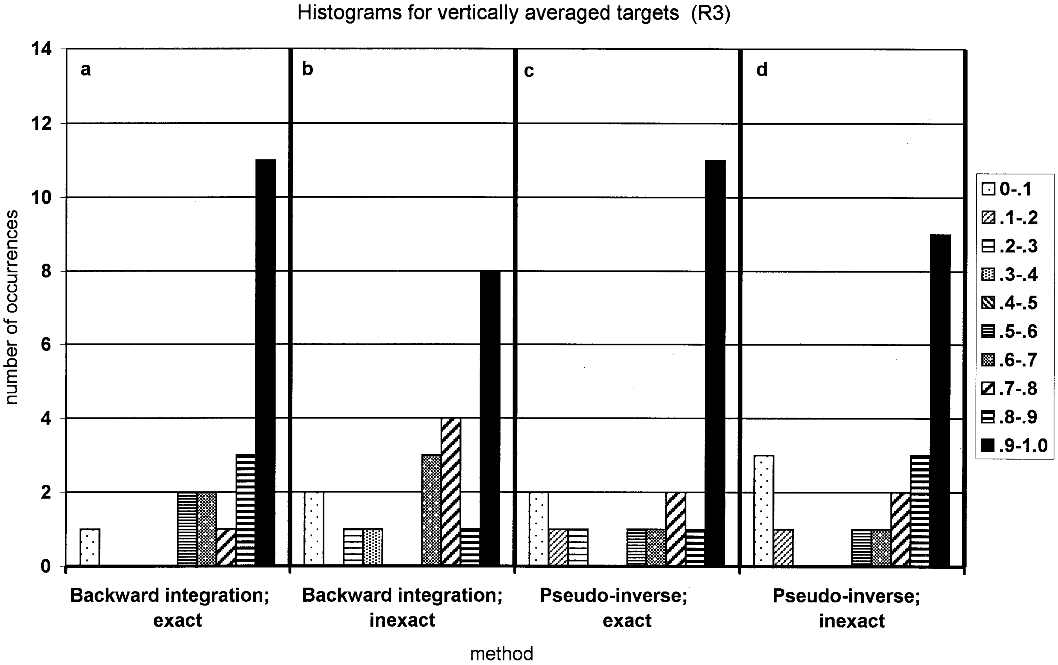

Fig. 8. Same as Fig. 7 but for the R3 analysis perturbations.

amplitude at 200 hPa than the growing SVs at initial time (Reynolds and Palmer, 1998; Reynolds, 1999).

These results are consistent with previous studies that found that problems arise during full inverse or approximate full inverse calculations (such as the backward integration technique) when uncertainty is added to the forecast error. These problems occur due to the erroneous projection of this noise or uncertainty onto the decaying SVs, which grow when integrated backward in time. It should be noted that different types of "noise" may have a smaller projection onto the trailing singular vectors, and therefore this effect may be somewhat mitigated. Reynolds (1999) found that perturbations constructed to have energy spectra similar to the analysis error estimates of Daley and Mayer (1986) have backward-in-time growth similar to that of the OI-3dVAR differences used here. If the uncertainty fields are constructed to have the most energy at either the largest or smallest spatial scales, then the backward-in-time growth of this uncertainty is smaller. However, unless unrealistically large negative dissipation is used, this noise still

Table 1. *Average normalized ranks, and % forecast error reduction (in parentheses), for 1-level targets (R10 analysis perturbations)*

	All targets	200-hPa targets only	500-hPa targets only	800-hPa targets only
exact pseudo-inverse	0.82 (15%) 20 targets	0.70 (6%) 2 targets	0.84 (16%) 18 targets	0 targets
inexact pseudo-inverse	0.87 (18%) 20 targets	0.86 (19%) 2 targets	0.87 (18%) 18 targets	0 targets
exact backward integration	0.79 (14%) 20 targets	0.40 (1%) 4 targets	0.89 (18%) 16 targets	0 targets
inexact backward integration	0.57 (7%) 20 targets	0.42 (2%) 14 targets	0.90 (17%) 6 targets	0 targets

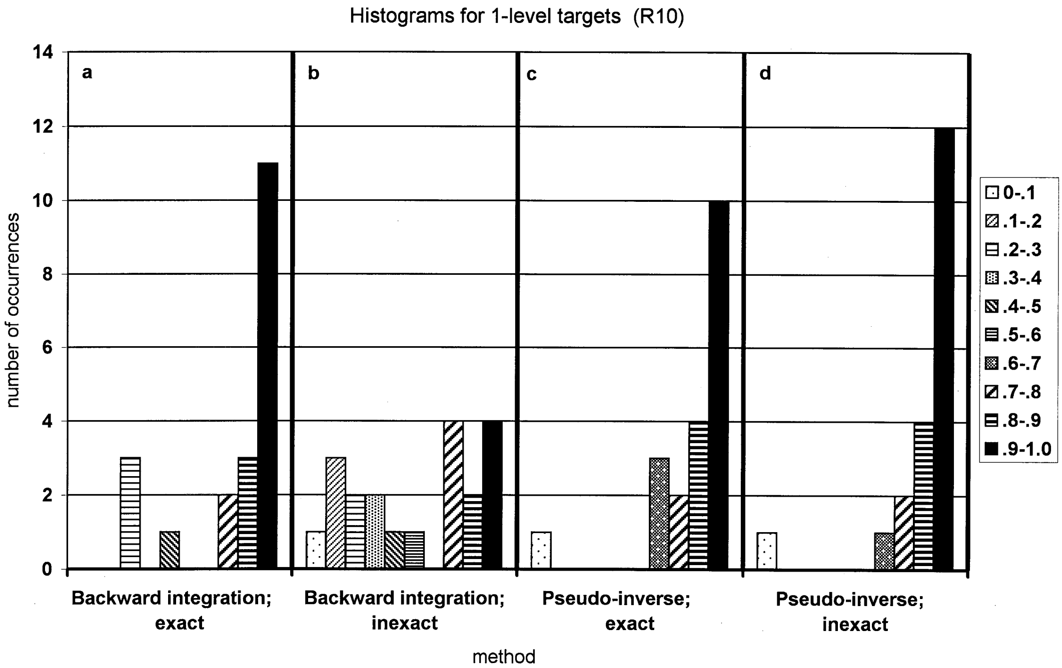

Fig. 9. Same as Fig. 7 but for the 1-level targets and R10 analysis perturbations.

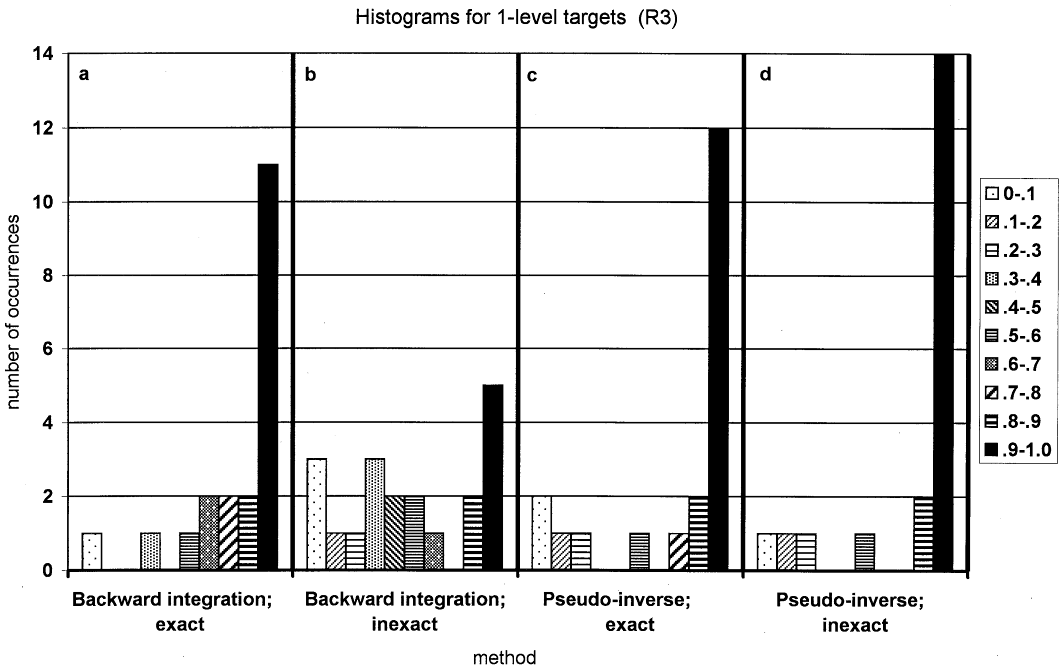

Fig. 10. Same as Fig. 7 but for the 1-level targets and R3 analysis perturbations.

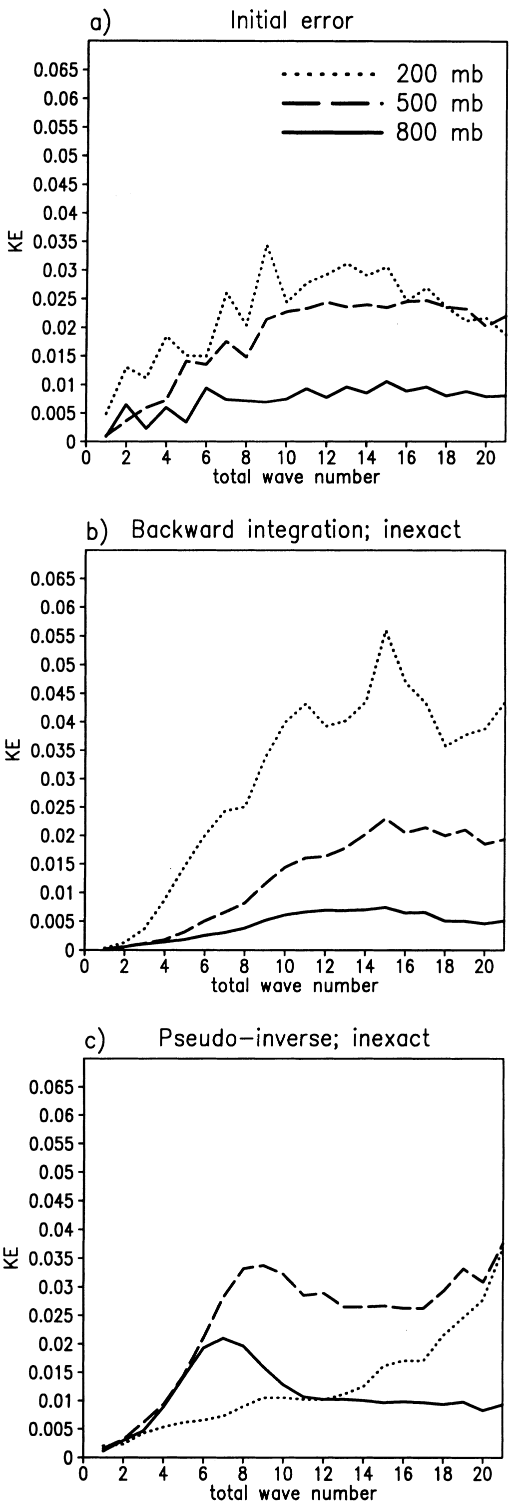

amplifies during the backward integration. Indeed, for these results to hold true, there need not be anything "special" about the characteristics of this noise or uncertainty other than that the projection of the uncertainty onto the final-time singular vectors is relatively flat when compared to the slope of the singular values. The problem is less significant with pseudo-inverse and adjoint techniques because these techniques focus on only the fastest growing component of the error, which has a much larger signal to noise ratio at final time than do the decaying components of the error.

All target rankings presented so far have been based on exact forecast error corrections. That is, the corrected forecast error is measured as the difference between the corrected forecast and the true atmospheric state (t_{48}). But the "inexact" techniques are designed to minimize the difference between the forecast error and the observed, or inexact, analysis (a_{48}). If this inexact analysis is used as verification, the results are slightly different. Fig. 12 is the same as Fig. 9, but a_{48} is used as verification rather than t_{48}. The pseudo-inverse results are similar when using either exact or inexact verification. The backward integration results are also similar, however, the exact backward integration targets appear somewhat worse, while the inexact backward integration targets appear somewhat improved. The inexact backward integration is still suboptimal when compared to the other methods. The small impact of using the inexact analyses rather than the true state for verification is not surprising given that the verifying analysis uncertainty, δ, is small compared with the forecast error. On average, the mean square geopotential height of δ over the verification region is 114 m^2, while the mean square geopotential height of $e_{\mathrm{f\,exact}}$ is 407 m^2.

The exact backward integration technique is effective even though this technique is designed to find both growing and decaying parts of the initial perturbation that evolves into the final-time per-

Fig. 11. The kinetic energy as a function of total wave number for (a) the initial analysis error, (b) the initial-time perturbation based on the inexact backward integration and (c) the initial-time perturbation based on the inexact pseudo-inverse. The energy at 200 hPa is denoted by the dotted line, 500 hPa by the dashed line and 800 hPa by the solid line. All fields have been normalized to have a total kinetic energy of 1.

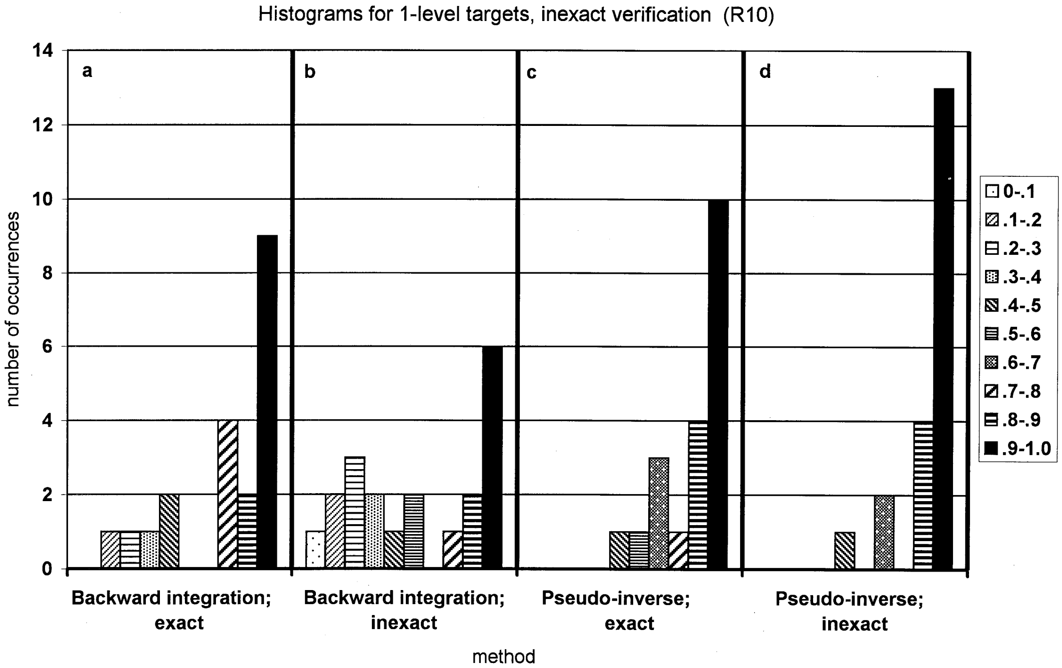

Fig. 12. Same as Fig. 9 but using inexact verification (i.e., using a_{48} as truth rather than t_{48}).

turbation over the verification region. The backward integration targets should then coincide with the largest part of this initial perturbation, which is not necessarily the fastest growing part. (In contrast, the pseudo-inverse is designed to find the fastest growing part of the initial perturbation.) There are some possible explanations for the effectiveness of the backward integration targets. First of all, the growing and decaying parts of the analysis error are often collocated (work not shown indicates that growing and decaying SVs are often horizontally collocated in the QG model). Also, the largest initial error may also be a fast growing error. Targets based on the PV maxima in the initial analysis errors (obviously not available in an operational setting) using the R10 analysis perturbations are shown in Fig. 13. The fact that 9 out of the 20 targets have a rank greater than 0.9 indicates that the largest initial errors are often fast growing. Note that this may be due to the simplicity and coarse resolution of the model used here, and may not be true for more complex systems.

Upon closer comparison of the different targeting methods, it is indeed found that the backward integration targets with high ranks are often those that are collocated either with the pseudo-inverse or maximum initial error targets. For example, for the one-level analysis perturbations (R10 or R3), the exact backward integration targets have ranks $\geqslant 0.9$ in 11 out of the 20 cases. Of those 11 high ranking backward integration targets, eight are collocated with (or within one gridpoint of) either the pseudo-inverse or maximum initial error targets. Similar ratios exist for the column analysis perturbations. On the other hand, lower ranking (< 0.9) backward integration targets are collocated with either the pseudo-inverse or maximum initial error targets in only two out of 12 cases. It should be noted that there are three cases when the backward integration targets have high ranks yet are not collocated with either the pseudo-inverse or maximum initial-error targets.

4.3. *Composite results for real-time methods*

All results presented so far require knowledge of the forecast error, and thus are only applicable in a post-time setting. Techniques available in real-time are now briefly examined. For opera-

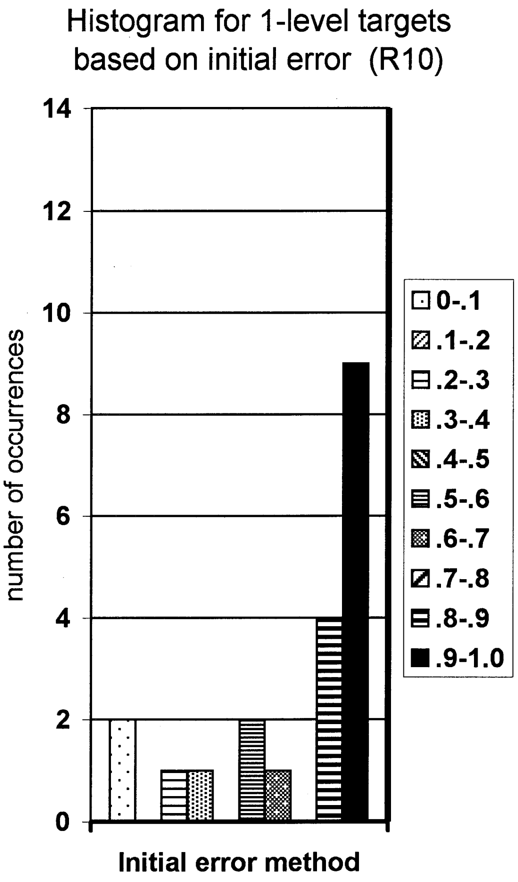

Fig. 13. Histogram of the rankings of the R10 1-level analysis perturbations for targets based on the largest initial-time analysis error.

tional purposes these methods would still need to be altered to allow for the lead time necessary for the deployment of the additional observing platforms.

Pu et al. (1998, 1999) suggest using the difference between two ensemble members as an estimate for the true forecast error. Ensemble backward integration methods have been used at NCEP in support of adaptive observation field programs. As we shall see, the effectiveness of the ensemble-based backward integration targets is very sensitive to the characteristics of the initial ensemble difference e_{p}. The ensembles are simulated by perturbing the initial conditions of the control forecast and making a third nonlinear forecast, as shown in Fig. 1. The difference between these two forecasts ($e_{\mathrm{f\,ensemble}}$) is used instead of the exact or inexact forecast errors ($e_{\mathrm{f\,exact}}$ or $e_{\mathrm{f\,inexact}}$) in determining the targets. As described in Subsection 3.2, three types of ensemble perturbations, $e_{\mathrm{p\,corr}}$, $e_{\mathrm{p\,lv}}$ and $e_{\mathrm{p\,uncorr}}$, are considered.

Results for the ensemble backward integration targets based on the 1-level R10 analysis perturbations are shown in Fig. 14. The correlated ensemble backward integration targets (Fig. 14a) are significantly better than the inexact backward integration targets (Fig. 9b). This is true even though the ensemble forecast difference is much worse than the inexact forecast error as an approximation to the exact forecast error. [The mean square geopotential height over the verification region for $e_{\mathrm{f\,exact}} - e_{\mathrm{f\,ensemble}}$ is 408 m^2, compared to 114 m^2 for δ ($e_{\mathrm{f\,exact}} - e_{\mathrm{f\,inexact}}$).] The Lyapunov ensemble backward integration targets (Fig. 14b) are likewise significantly better than the inexact backward integration targets. On the other hand, the uncorrelated ensemble backward integration targets (Fig. 14c) appear to be only marginally more optimal than targets chosen at random. This is true even though the uncorrelated ensemble backward integration targets do occur far less often at the 200-hPa level (in 3 out of 20 cases) than the inexact backward integration targets (14 out of the 20 cases). Several other different types of uncorrelated ensemble perturbations (e.g., perturbations with more energy at smaller or larger wave numbers, and perturbations that are barotropic) all gave similar results to those shown here.

Results for the ensemble pseudo-inverse targets are shown in Fig. 15. Comparison with Figs. 9c,d indicates that the ensemble pseudo-inverse targets are comparable with the exact and inexact pseudo-inverse targets. Similar results are obtained by using the SVs only (assuming a flat projection of the initial-time error field onto the SVs, not shown). This confirms that the exact projections of the initial-time error onto the SVs are not needed in order to obtain good targets.

Results based on the vertically averaged targets and column perturbations are qualitatively similar (Table 2). However, except for the uncorrelated ensemble backward integration method, the vertically averaged targets have smaller ranks than the 1-level targets.

Upon consideration of the nature of these techniques, the results are not surprising. An ensemble forecast difference is dominated by the growing SVs. Since the projection of the forecast difference

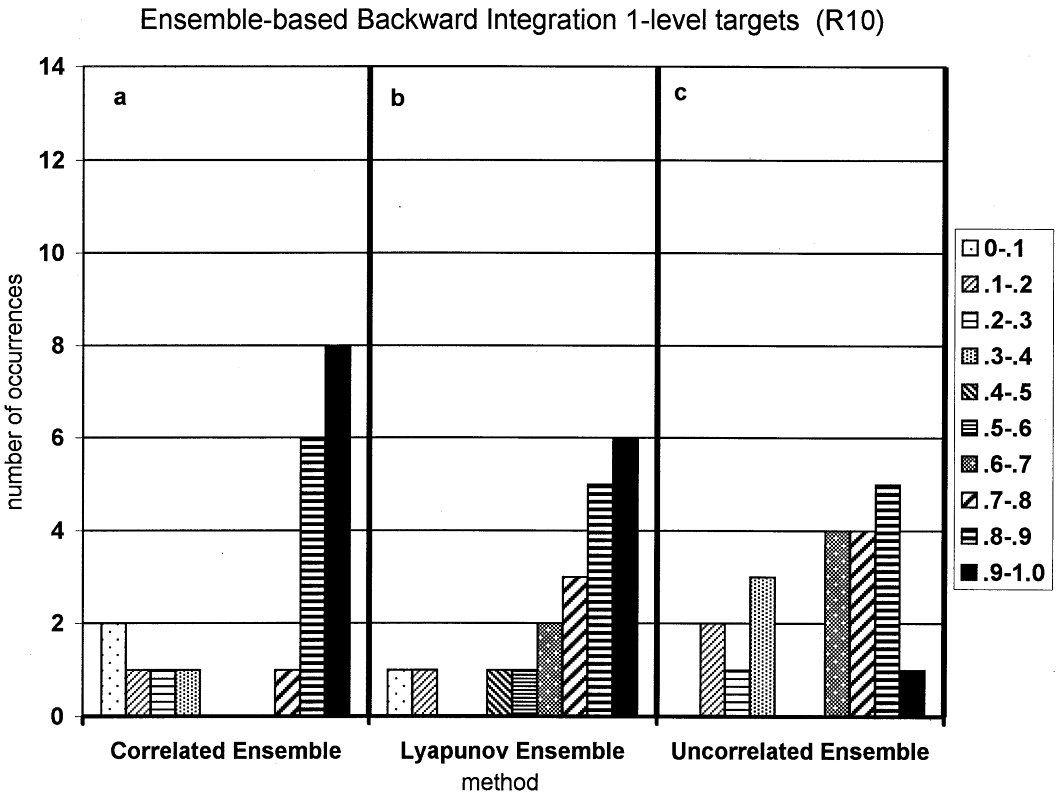

Fig. 14. Histograms of the rankings of the R10 1-level analysis perturbation backward integration targets based on (a) correlated ensemble perturbations, (b) Lyapunov ensemble perturbations, and (c) the uncorrelated ensemble perturbations.

onto the decaying SVs is very small, the spurious 200-hPa targets produced by the inexact backward integration do not occur when using the ensemble backward integration. But in order for the ensemble backward integration targets to be effective, the initial ensemble perturbation must either resemble the true analysis error (as in $e_{\mathrm{p\,corr}}$) or be dominated by growing errors (as in $e_{\mathrm{p\,lv}}$). (Of course, if the initial ensemble perturbation were composed of the 3 leading singular vectors, then the backward integration and pseudo-inverse results would be very similar.) In contrast, the location of the pseudo-inverse target is constrained to only a few possible locations because the initial-time perturbation is composed of some linear combination of only three singular vectors. It should be noted that the ensemble pseudo-inverse targets usually have higher average ranks than the ensemble backward integration targets (Table 2). The exception is for the 1-level vertically averaged Lyapunov ensemble targets, where the average rank of the backward integration targets (0.75) is slightly larger than the average rank of the pseudo-inverse targets (0.72). The ensemble pseudo-inverse targets are better than the ensemble backward integration targets in more cases than vice versa (Table 2).

Targets based on features of the initial-time analysis have also been examined (not shown). Targets based on the maxima in the horizontal PV gradients in the analyses at the 200-hPa and 500-hPa levels are comparable to picking targets at random. Targets based on the maximum PV in the analyses are likewise ineffective.

5. Summary and conclusions

A T21L3 QG model has been used to examine the optimality of different targeting methods in an

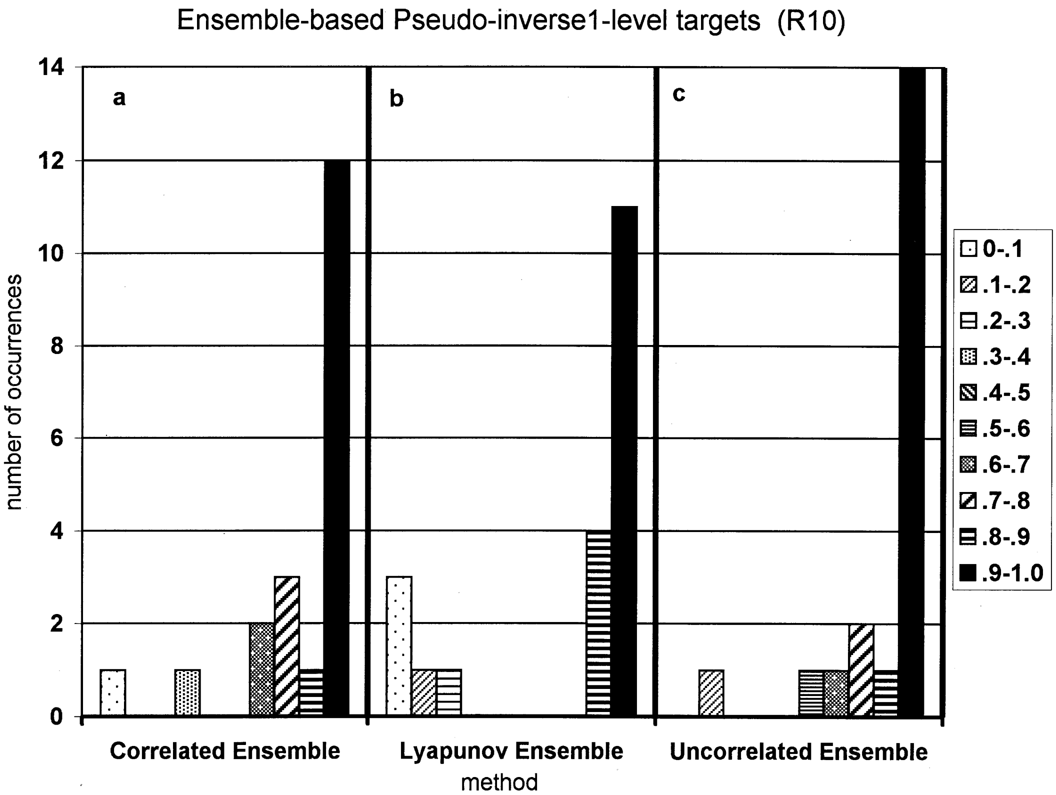

Fig. 15. Same as Fig. 14 but for the pseudo-inverse targets.

Table 2. *First two columns show the averaged normalized ranks, and % forecast error reduction (in parentheses), for the ensemble pseudo-inverse and ensemble backward integration targets; last two columns show the number of times the ensemble pseudo-inverse target is better than the ensemble backward integration target (or vice versa); targets are considered equivalent if ranks are within 3% of each other; results are for R10 analysis perturbations*

	Ensemble pseudo-inverse average rank	Ensemble backward integration average rank	# of cases pseudo-inverse > backward int.	# of cases backward int > pseudo-inverse
vertically-averaged target: correlated ensemble	0.77 (23%)	0.67 (17%)	13	6
vertically-averaged target: Lyapunov ensemble	0.77 (25%)	0.66 (18%)	12	5
vertically-averaged target: uncorrelated ensemble	0.78 (23%)	0.63 (19%)	13	4
1-level target: correlated ensemble	0.82 (15%)	0.71 (10%)	11	7
1-level target: Lyapunov ensemble	0.72 (13%)	0.75 (11%)	9	5
1-level target: uncorrelated ensemble	0.85 (15%)	0.62 (6%)	17	2

idealized setting. Both the backward integration and pseudo-inverse techniques provide targets that are considerably better than targets chosen at random when computed in an a posteriori setting based on the exact forecast error. When uncertainty in the verifying analysis is added to the forecast error, the effectiveness of the backward integration targets is considerably degraded. This is due to the erroneous projection of the inexact forecast error onto the decaying SVs, which amplify during the backward integration. This results in spurious targets that, for this model, are often found at the 200-hPa level. If an ensemble-forecast difference is used instead of the inexact forecast error, the problem of spurious targets at the upper levels is mitigated considerably. However, the usefulness of the ensemble backward integration targets depends on the character of the initial ensemble perturbation. Ensemble perturbations should either resemble the true analysis errors or be dominated by the growing part of the initial error (or some combination there of). If this occurs, then the ensemble backward integration targets are better than the inexact backward integration targets, and in some cases, comparable to the pseudo-inverse targets. The pseudo-inverse targets are not particularly sensitive to the initial-time perturbation. The exact, inexact and ensemble pseudo-inverse targets all give good results.

Pu et al. (1998) point out the differences between backward integration and adjoint methods and suggest using both methods for targeting. In the present study, if we consider only those cases in which the backward integration and pseudo-inverse targets are collocated, then the average ranking is considerably higher than the 20-case average, suggesting that this collocation might be used as an indicator of target effectiveness. A larger sample size is needed to confirm this relationship. For this model, targets based on features of the analysis, such as the largest 200-hPa or 500-hPa PV horizontal gradients, are not particularly effective.

The exact backward integration targets can be effective even though the technique is designed to find both the growing and decaying parts of the initial analysis error that evolve into the forecast error over the verification region. (In contrast, the pseudo-inverse is designed to find only the fastest growing part of that initial perturbation.) One reason for this may be the fact that the growing and decaying SVs in this system are often collocated. A second reason may be the fact that large analysis errors are often also fast growing errors in this system. Most of the high-ranking backward integration targets are indeed collocated with either the pseudo-inverse or maximum initial error targets. It should be noted, however, that the growing and decaying SVs might not be collocated in a more complex system. Likewise, the fact that the large analysis errors often serve as good targets may be an artifact of the characteristics of the prescribed initial perturbations and the simplicity of the model. Results using an operational system could be quite different.

There are other limitations of this study, primarily owing to its simplicity, that should be kept in mind when interpreting the results. Firstly, the model has a very coarse horizontal and vertical resolution, and the impact of the simple analysis perturbations may be a strong function of this coarse resolution. Secondly, the perfect model assumption has significant implications, not only in the context of errors in the operational forecast model, but also in differences between the nonlinear operational forecast model and the tangent linear model. In an operational setting, the tangent equations that are used during the backward integration and to calculate the SVs are often based on a dry model that is considerably simpler than the model used to produce the nonlinear control and ensemble forecasts. In this situation it is conceivable that the nonlinear ensemble differences *may* have a significant projection onto the decaying SVs of the simpler linear system. If this is true, then the effectiveness of the ensemble backward integration targets may be overly optimistic. Likewise, though dry SVs can often account for a large portion of the operational forecast error, moist processes may be very important in some cases (although there is evidence that moist and dry SVs are often collocated, even if their structures and growth rates are considerably different, Ehrendorfer et al., 1999; Gelaro et al., 1999).

The effectiveness of backward-integration ensemble-based targets will depend on how much information about the true analysis error and/or the growing component of that error can be extracted from the ensemble perturbations. For operational systems, a better way to extract this information may be to use the variance between all the ensemble members rather than simply picking any two ensemble members at random. The ensemble perturbations currently used at operational centers probably do not bear a significant resemblance to the true analysis errors on any given day. At ECMWF, for example, a better estimate of the analysis error covariance is needed in order to provide the optimal metric for the singular vectors at initial time (Ehrendorfer and Tribbia, 1997; Barkmeijer et al., 1998). As for the bred modes used at NCEP, Iyengar et al. (1996) estimate that five bred modes explain only roughly 4% of the estimated analysis error. This may be due, at least in part, to the fact that the rescaling of the perturbations does not adequately account for the scale-dependent effects of data assimilation.

It should be noted, however, that if the ensemble perturbations are composed primarily of growing components, the backward integration targets based on them should be more robust. This is evident in the relatively good results obtained with the Lyapunov ensemble backward integration targets. In practice, choosing the 2 ensemble members that exhibit the largest differences in the target region at verification time should also help insure that significant growing components of the perturbation are contained within the ensemble difference. Should the initial ensemble perturbation be constructed from the first three singular vectors, then the backward integration and pseudo-inverse results would be very similar.

In this simple study, the issue of irregular data density is not addressed at all. In an operational setting, the regional differences in the conventional data network should be taken into account in the targeting method. With the pseudo-inverse, this might be accomplished by using an appropriate initial-time error variance or covariance norm. Work on development of an analysis error variance norm is currently underway at ECMWF and the Naval Research Laboratory. When the backward integration technique is used at NCEP, this issue is partly addressed when the bred modes used for the ensemble perturbations are scaled to reflect regional differences in estimated analysis error variance (Toth and Kalnay, 1997).

The target gridpoints are determined from the initial-time perturbation fields in a very simplistic manner. There are certainly more sophisticated, and perhaps more effective, ways to pick the targets, particularly when there is more than one local maximum in the initial-time perturbation field. Also, the analysis perturbation structures are highly idealized. Vautard et al. (1999) have shown that the effectiveness of targeting is very sensitive to the way in which the target is sampled. They find that just sampling the extrema of a sensitivity pattern is far less effective for improving forecasts then a more complete sampling of the sensitivity pattern. It would be interesting to try more sophisticated analysis perturbations whose shape and vertical structure would be based on the initial-time perturbation. Addressing these issues may be a fruitful avenue for future work

6. Acknowledgments:

The 3-level quasi-geostrophic model was developed by Franco Molteni. The authors would like to thank Ron Errico, Eugenia Kalnay, and ZhaoXia Pu for comments on an earlier draft of this paper. L. A. Smith suggested the experiment shown in Fig. 12. Support of the sponsor, Office of Naval Research, and the program manager, Naval Research Laboratory, program element 0601153N, program number BE-033-0345, is gratefully acknowledged.

REFERENCES

Barkmeijer, J., van Gijzen, M. and Bouttier, F. 1998. Singular vectors and the analysis error covariance metric. *Q. J. R. Meteorol. Soc.* **124**, 1695–1713.

Barkmeijer, J. 1992. Local error growth in a barotropic model. *Tellus* **44A**, 314–323.

Bergot, T., Hello, G., Joly, A. and Malardel, S. 1999. Adaptive observations: a feasibility study. *Mon. Wea. Rev.* **127**, 743–765.

Buizza, R., Gelaro, R., Molteni, F. and Palmer, T. N. 1997. The impact of increased resolution on predictability studies with singular vectors. *Q. J. R. Meteorol. Soc.* **123**, 1007–1033.

Buizza, R. 1994. Localization of optimal perturbations using a projection operator. *Q. J. R. Meteorol. Soc.* **120**, 1647–1681.

Daley, R. and Mayer, T. 1986. Estimates of global analysis error from the Global Weather Experiment Observational Network. *Mon. Wea. Rev.* **114**, 1642–1653.

Ehrendorfer, M., Errico, R. M. and Raeder, K. D. 1999. Singular-vector perturbation growth in a primitive-equation model with moist physics. *J. Atmos. Sci.* **56**, 1627–1648.

Ehrendorfer, M. and Tribbia, J. J. 1997. Optimal prediction of forecast difference covariances through singular vectors. *J. Atmos. Sci.* **54**, 286–313.

Ehrendorfer, M. and Errico, R. M. 1995. Mesoscale predictability and the spectrum of optimal perturbations. *J. Atmos. Sci.* **52**, 3475–3500.

Gelaro, R., Langland, R. H., Rohaly, G. D. and Rosmond, T. E. 1999. An assessment of the singular vector approach to targeted observations using the FASTEX data set. *Q. J. R. Meteorol. Soc.* **125**, 3299–3328.

Gelaro, R., Buizza, R., Palmer, T. N. and Klinker, E. 1998. Sensitivity analysis of forecast errors and the construction of optimal perturbations using singular vectors. *J. Atmos. Sci.* **55**, 1012–1037.

Iyengar, G., Toth, Z., Kalnay, E. and Woolen, J. S. 1996. *Are the bred vectors representative of analysis errors*? Preprints: 11th AMS Conference on Numerical weather prediction. Norfolk, VA, American Meteorological Society, J64–J66.

Joly, A., Browning, D. A. Bessemoulin, P., Cammas, J.-P., Caniaux, G., Chalon, J.-P., Clough, S. A., Dirks, R., Emanuel, K. A., Eymard, L., LaLaurette, F., Gall, R., Hewson, T. D., Hildebrand, P. H., Jorgensen, D., Langland, R. H., Lemaitre, Y., Mascart, P., Moore, J. A., Perrson, P. O. G., Roux, F., Shapiro, M. A., Snyder, C., Toth, Z. and Wakimoto, R. M. 1999. Overview of the field phase of the Fronts and Atlantic Storm-Track Experiment (FASTEX) project. *Q. J. R. Meteorol. Soc.* **125**, 3131–3164.

Langland, R. H., Gelaro, R., Rohaly, G. D. and Shapiro, M. A. 1999a. Targeted observations in FASTEX: adjoint-based targeting procedures and data impact experiments in IOPs-17 and 18. *Q. J. R. Meteorol. Soc.* **125**, 3241–3270.

Langland, R. H., Toth, Z., Gelaro, R., Shapiro, M. A., Szunyogh, I., Majumdar, S. J., Morss, R., Rohaly, G. D., Velden, C., Bond, N. and Bishop, C. H. 1999b. The North Pacific Experiment — 1998 (NORPEX — 98): targeted observations for improved North American weather forecasts. *Bull. Amer. Meteor. Soc.* **80**, 1363–1384.

Marshall, J. and Molteni, F. 1993. Toward a dynamical understanding of planetary-scale flow regimes. *J. Atmos. Sci.* **50**, 1792–1818.

Molteni, F. and Palmer, T. N. 1993. Predictability and finite-time instability of the northern winter circulation. *Q. J. R. Meteorol. Soc.* **119**. 269–298.

Palmer, T. N., Gelaro, R., Barkmeijer, J. and Buizza, R. 1998. Singular vectors, metrics and adaptive observations. *J. Atmos. Sci.* **55**, 633–653.

Pu, Z.-X. and Kalnay, E. 1999. Targeting observations with the quasi-inverse linear and adjoint NCEP global models: Performance during FASTEX. *Q. J. R. Meteorol. Soc.* **125**, 3329–3338.

Pu, Z.-X., Lord, S. and Kalnay, E. 1998. Forecast sensitivity with dropwindsonde data and targeted observations. *Tellus* **50A**, 391–410.

Pu, Z.-X., Kalnay, E., Sela, J. and Szunyogh, I. 1997. Sensitivity of forecast errors to initial conditions with a quasi-inverse linear method. *Mon. Wea. Rev.* **125**, 2479–2503..

Rabier, F., Klinker, E., Courtier, P. and Hollingsworth, A. 1996. Sensitivity of forecast errors to initial conditions. *Q. J. R. Meteorol. Soc.* **122**, 121–150.

Reynolds, C. A. 1999. A comparison of analysis and forecast correction techniques: impact of negative dissipation. *Mon. Wea. Rev.* **127**, 2576–2596.

Reynolds, C. A. and Palmer, T. N. 1998. The impact of decaying singular vectors on analysis and forecast correction. *J. Atmos. Sci.* **55**, 3005–3023.

Reynolds, C. A. and Errico, R. 1999. Convergence of singular vectors towards Lyapunov vectors. *Mon. Wea. Rev.* **127**, 2309–2323.

Szunyogh, I., Toth, Z., Emanuel, K. A., Bishop, C. H., Snyder, C., Morss, R., Wollen, J. and Marchok, T. 1999. Ensemble-based targeting experiments during FASTEX: the impact of dropsonde data from the LEAR jet. *Q. J. R. Meteorol. Soc.* **125**, 3189–3218.

Toth, Z, and Kalnay, E. 1997. Ensemble forecasting at NCEP and the breeding method. *Mon. Wea. Rev.* **125**, 3297–3319.

Wang, Z., Droegemeier, K. K., White, L. and Navon, I. M. 1997. Application of a new adjoint newton algorithm to the 3D ARPS storm-scale model using simulated data. *Mon. Wea. Rev.* ***125***, 2460–2478.