Tellus (2001), 53A, 317–332

TELLUS
ISSN 0280–6495

The spectra of singular values in a regional model

By RONALD M. ERRICO[1]*, MARTIN EHRENDORFER[2] and KEVIN RAEDER[1], [1]*National Center for Atmospheric Research, USA*†; [2]*Institute for Meteorology and Geophysics, University of Vienna, Austria*

(Manuscript received 28 March 2000; in final form 5 December 2000)

ABSTRACT

Large portions of the spectra of singular values are determined for both moist and dry versions of a tangent linear, regional model for 4 different synoptic cases. Norms considered include the usual energy norm and versions of a norm measuring only the energy in some set of rotational mode perturbations. At most, only a few percent of the singular vectors possible with any of the norms are growing ones. Inclusion of moist physics in the tangent linear model greatly affects the leading singular vectors but does not increase the number of growing singular vectors much. Most singular vectors are damping ones, and therefore random perturbations drawn from a white-noise distribution will likely damp during the 24-h forecast periods considered. Only a few singular vectors are required to explain a significant portion of the final-time variance of such perturbations, however, because the leading singular values are so large compared with the rest. The truncated rotational mode norm is shown to be very useful for investigating these properties.

1. Introduction

Singular vectors (SVs) are determined for many types of applications in dynamic meteorology. These applications include stability analysis (Farrell and Ioannou, 1996a, b), ensemble forecasting (Buizza, 1994), observation targeting (Gelaro et al., 1999), forecast and analysis error estimation (Ehrendorfer and Tribbia, 1997; Gelaro et al., 1997), and data assimilation (Cohn and Todling, 1996). For many problems, SVs determined using appropriate norms can be shown to yield optimal solutions (Farrell, 1989; Ehrendorfer and Tribbia, 1997). For example, the SVs appropriately determined for a linearized model define an orthogonal set of perturbation structures ordered from fastest growing (or leading) to fastest decaying (or trailing) structures when growth is measured using some specified norm over some specified period (see the review by Ehrendorfer (1997) for some other discussion and examples). The leading SV is an optimal perturbation in the specific sense that any other perturbation will experience no greater net growth during that period in the linearized model.

Moist diabatic processes have long been known to have profound effects on atmospheric evolution in general and rapid cyclogenesis in particular (Mullen and Baumhefner, 1989). Without the availability of suitable linearizations of moist diabatic processes, however, SV determinations had been limited to dry versions of linearized models (Errico, 1997). Recently, useful linearizations of such processes have been developed and documented (Mahfouf, 1999; Errico and Raeder, 1999). SV determinations incorporating those processes have demonstrated a sometimes profound impact on both the growth rates and structures of leading SVs (Ehrendorfer et al., 1999; Barkmeijer et al., 2001).

* Corresponding author address:
Ronald M. Errico, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307-3000.
e-mail: rmerrico@ucar.edu.
† The National Center for Atmospheric Research is sponsored by the US National Science Foundation.

Usually, only a small set of leading SVs are determined, because for most applications only the few fastest growing SVs are required and determination of a larger set can be computationally prohibitive. One exception is Reynolds and Palmer (1998), who determine an entire set of SVs, including the trailing ones, for a low-resolution quasi-geostrophic model. In most studies with more realistic models, however, only 1–50 SVs are determined out of a possible set of 10^5 or more (Gelaro et al., 1999; Palmer et al., 1998, and references therein).

Knowledge of even some non-leading SVs would be valuable. How fast do the growth rates of successive SVs drop off? In particular, how many growing SVs are there in some typical synoptic cases? What are the characteristics of the more slowly growing SVs? What are the differences between subspaces of leading SVs produced using different norms? Answers to these questions reveal aspects of the likely behavior of random initial perturbations (Farrell, 1990). In particular, under conditions described in Section 5, they can be used to establish bounds on the growth of such perturbations, and thus can be used to characterize predictability of a synoptic case (Ehrendorfer and Errico, 1995).

The most commonly used norm is known as the total energy norm (Palmer et al., 1998; Talagrand, 1981; Errico, 2000a). A norm based on it but with other desirable properties, called the rotational mode norm, was recently introduced by Errico (2000a), but he examined only the leading few SVs. How answers to the above questions are dependent on these norms is examined here. Also asked is how the results are affected by the truncation specified for the rotational mode norm and by the additional consideration of moist physics to the dry dynamics of the model. The explicit consideration of moisture by the norm itself (i.e., explicitly measuring moisture in its formulation) is not considered here (but see Ehrendorfer et al. (1999) for an example of this).

Some of these same questions were explored by Ehrendorfer and Errico (1995) but for a more limited set of cases and norms. Although they also used a norm based on rotational modes, it had to be restricted in vertical scale and modified so it no longer described the energy of the modes, for reasons described by them and by Ehrendorfer et al. (1999). For this new study, no restriction or modification of the rotational mode norm is required, larger sets of SVs are determined for each norm and synoptic case, and more synoptic cases are examined.

The SVs and their norms are described in Section 2, followed by a summary of the particular model and cases investigated in Section 3. Results are presented in Sections 4 and 5, the latter describing their statistical implications. In Section 6, the similarity between subspaces of SVs for pairs of norms are described. Conclusions are offered in Section 7.

2. Singular vectors and norms

In this study, attention is restricted to consideration of the space of a model's perturbed initial conditions (or some portion thereof, with any remaining portion set to zero) defined in some discrete way. A perturbation within that subspace is denoted as the vector $\boldsymbol{x}(0)$ and its linearized forecast at time t as $\boldsymbol{x}(t)=\boldsymbol{M}\boldsymbol{x}(0)$ where $\boldsymbol{M}$ is the resolvent of the linearized model at t. Only the problem posed by applying the same norm at both initial and final times is examined here. That norm is described as

$$\|\boldsymbol{x}\|^2=\boldsymbol{x}^{\mathrm{T}}\boldsymbol{N}\boldsymbol{x}, \tag{2.1}$$

where $\boldsymbol{N}$ is the positive definite matrix defining the particular norm and superscript T denotes a transpose.

The SVs may be determined as $\boldsymbol{x}_k(0)=\boldsymbol{N}^{-1/2}\boldsymbol{z}_k$ where $\boldsymbol{z}_k$ are the eigenvector solutions to

$$\boldsymbol{N}^{-1/2\mathrm{T}}\boldsymbol{M}^{\mathrm{T}}\boldsymbol{N}\boldsymbol{M}\boldsymbol{N}^{-1/2}\boldsymbol{z}=\lambda^2\boldsymbol{z}, \tag{2.2}$$

ordered by descending values of their corresponding eigenvalues λ^2. The $\boldsymbol{x}_k(0)$ are the right singular vectors of $\boldsymbol{M}$ with respect to the norm defined by $\boldsymbol{N}$, the λ are the corresponding singular values, and $\boldsymbol{x}_k(t)$ are the corresponding left singular vectors. In the meteorological literature $\boldsymbol{x}_k(0)$ are often simply called SVs and the $\boldsymbol{x}_k(t)$ are called evolved SVs. The SV associated with the largest λ is known as the leading SV and is denoted here as $\boldsymbol{x}_1(0)$. Distinct SVs are orthogonal with respect to the norm specified by $\boldsymbol{N}$ at both initial and final times.

Since, for large realistic models, the matrix $\boldsymbol{M}$ exists only as an operator in the form of a computer code, the SVs are most easily obtained using an iterative Lanczos algorithm (Strang, 1986; Grimes et al., 1994) that does not require an explicit representation of the model operator. Estimates for larger values of λ converge more rapidly. In the experiments here, for the number of iterations (I) of the Lanczos algorithm, estimates of approximately $I/2$ of the λ_k^2 have converged to within 0.001% of their true values when $I > 300$. Only such converged values are presented here.

The total energy (or E-) norm appropriate for the model to be used here is expressed by a diagonal N (Ehrendorfer et al., 1999) with elements $\Delta A/2$ times $\Delta\sigma$ applied to eastward and northward wind components (u, v, respectively), times $\Delta\sigma c_p/T_r$ for temperature (T), and times RT_r/p_{sr}^2 for surface pressure (p_s), where ΔA and $\Delta\sigma$ are the horizontal area and vertical depth (in terrain following σ coordinates) of a model grid box, c_p is the specific heat of air at constant pressure, R is the gas constant for dry air, $T_r = 270$ K is a spatially independent reference T, and $p_{sr} = 100$ kPa is a horizontally independent reference surface pressure. In the specific model used here, the $\Delta\sigma$ are the same at all model levels, and the ΔA will be approximated by a spatially constant value. Some properties of this norm are discussed in Talagrand (1981) and more completely for this model in Errico (2000b). The dimension (N) of the state vector measured by this norm is equal to the number of independent forecast values of u, v, T, p_s defined on the model grid.

The rotational mode norm used here is described in detail in Errico (2000a). It is that portion of the E-norm expressible in terms of rotational modes of a simple reference model derived from the model to be employed. These modes are geostrophic and describe a linearized form of potential vorticity. Energy ascribed to gravitational modes (including inertial-gravity waves) is thereby excluded from the norm. SVs determined using this norm have no projection on gravitational modes at the initial time, and any subsequent energy in these modes is excluded from the measure of growth produced at the final time. For this norm, N is a little less than 1/3 of that of the E-norm since portions of the state vector projecting onto gravitational modes are excluded.

A truncated form of the R-norm (labeled as TR) is primarily considered here. It differs from the full R-norm in that it excludes consideration of horizontal wavelengths less than 4 grid lengths in either direction and consideration of the 3 vertical modes with equivalent depths less than 10 m. It is designed so that, compared to either the E-norm or non-truncated R-norm, it can more reasonably be interpreted as a quadratic measure of the initial perturbation weighted by an assumed, geostrophic and spatially correlated, analysis error covariance (Errico, 2000a). For this norm, N is approximately 18% of that for the full R-norm.

3. Model and cases investigated

The model used is identical to the version of the mesoscale adjoint modeling system (denoted as MAMS2) described in Errico and Raeder (1999). Pertinent details here are that it uses the primitive equations defined on 10 σ-levels (in addition to the surface) within a regional domain with dry diabatic process including vertical and horizontal diffusion, bulk formulation of the planetary boundary layer, dry convective adjustment, and a prognostic surface temperature over land. Moist diabatic processes include convective (Moorthi and Suarez, 1992) and stratiform precipitation and effects of clouds on radiation at the surface. Its tangent linear (TLM) and adjoint versions (Errico, 1997) required for SV determination are exact except for some minor approximations applied for computational efficiency as described in Errico and Raeder (1999). See Errico et al. (1994) for full description of its parent model (MAMS1).

Lateral boundaries are modeled using a Davies and Turner (1977) relaxation scheme acting within 5 points of the grid edges. In the linearized version of the model, this boundary relaxation corresponds to damping of perturbations. This damping and the apparent exit of perturbations from the domain due to outflow advection, with no corresponding inflow of perturbations, must act to diminish the number of possible growing SVs compared with an unbounded or periodic domain.

Initial conditions for the nonlinear reference forecasts are obtained from interpolations of global analyses produced operationally at the European Center for Medium-Range Weather Forecasts and archived at the National Center for Atmospheric Research. These archived data sets are at a relatively low resolution, defined on an approximately $3^\circ \times 3^\circ$ horizontal grid on

7 pressure levels in the vertical. A nonlinear vertical mode initialization scheme (Bourke and McGregor, 1983) is applied to the external and first internal modes of the interpolated fields. The fields used to prescribe the lateral boundaries are provided by similarly interpolated analyses every 12 h, with linear interpolations performed in time within each 12-h synoptic period.

Four synoptic cases are examined. All cover 24-h periods. The basic states about which the model is linearized are obtained from model forecasts produced with moisture and all the physics. The two winter case domains and grids are identical as are (independently) those for the two summer cases. The model top is at 1 kPa.

Case W1 begins 0000 UTC 14 February 1982. It is dominated by an explosively deepening cyclone off the Atlantic Coast of North America. The grid spacing is 120 km and the domain (but not the synoptic situation) is shown in Fig. 1. The numbers of degrees of freedom for the vector measured by the E-norm, R-norm, and TR-norm are 79 354, 24 940, and 4830, respectively. This is the same case whose leading SVs are described in Ehrendorfer et al. (1999), Errico and Raeder (1999), and Errico (2000a, b).

Case W2 begins 0000 UTC 13 January 1988. Its initial 50 kPa geopotential height and sea-level pressure (p_{sl}) fields as well as forecast p_{sl} and accumulated precipitation fields are shown in Fig. 1. Initially, the case is dominated by a surface high pressure centered near New York and a short-wave, low amplitude, 50 kPa trough over the Mississippi River. Although there was minimal surface cyclone development associated with the short wave in reality, the simulation does produce a weak oceanic cyclone ahead of it. The simulated precipitation is predominantly stratiform in the northeast but convective off the southern US coast. This case was chosen because it is unlike

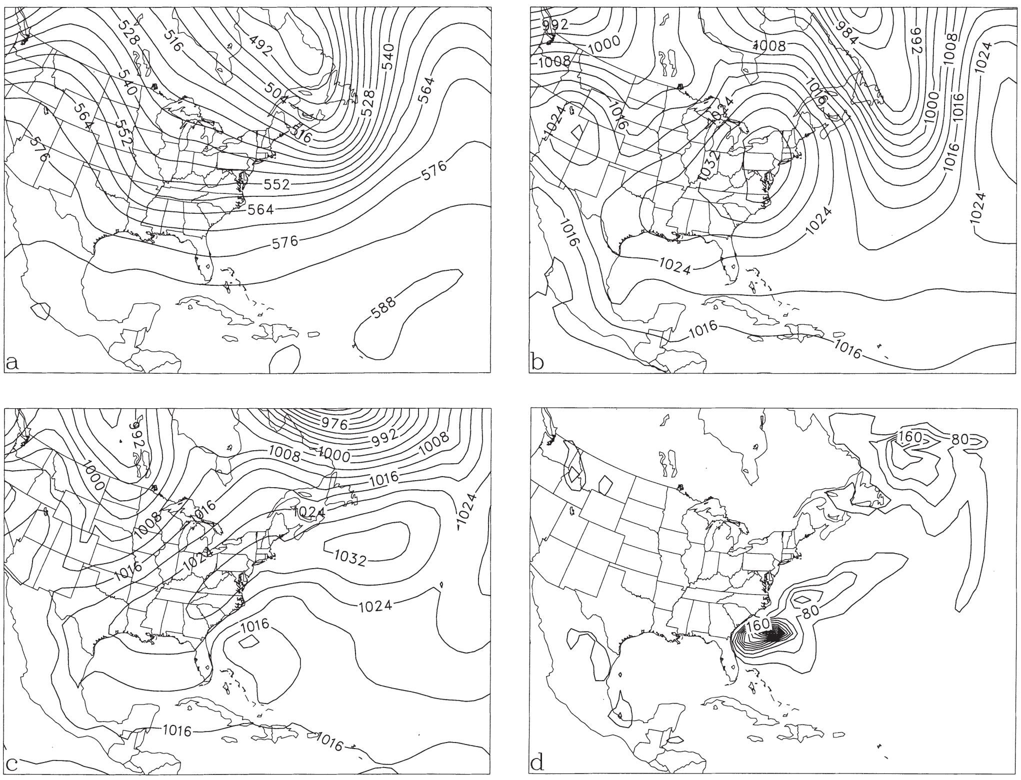

Fig. 1. Initial fields of (a) 50 kPa geopotential height z and (b) sea-level pressure p_{sl}, and final fields of (c) p_{sl} and (d) accumulated precipitation R for case W2. Contour intervals are 6 dam for z, 0.4 kPa for p_{sl}, and 0.4 mm for R.

W1 in that there is much weaker cyclogenesis both in reality and in the simulation.

Case S1 begins 1200 UTC 16 June 1989. It is dominated by two regions of precipitation: one over Saskatchewan and the other over the southern Appalachian Mountains. The grid spacing is 80 km and the domain (but not the synoptic situation) is shown in Fig. 2. The numbers of degrees of freedom for the vector measured by the E-norm and TR-norm are 65 080 and 3969, respectively. This is the same case whose leading SVs are described in Ehrendorfer et al. (1999) and Errico and Raeder (1999).

Case S2 begins 0000 UTC 3 July 1988. Its initial 50 kPa geopotential height and p_{sl} fields as well as forecast p_{sl} and accumulated precipitation fields are shown in Fig. 2. The case is dominated by a 50 kPa ridge that builds over most of the US and two very weak surface cyclones that undergo only weak development over the Rocky Mountains and Southern Canadian Plains. Most of the precipitation is convective, except for a small region of stratiform precipitation over Colorado. This case was chosen because, unlike in S1, most of the domain is dry.

4. Results

Results for all the cases and norms investigated are summarized in Table 1. Indicated there is the case number and label, the physics and norm used, the size N of the space measured by the norm, the number I of iterations of the Lanczos algorithm performed, the number I_c of converged λ^2 using the criteria described in Section 2, the number N_g of λ_k with values greater than 1, λ_1^2, $\lambda_{I_c}^2$, and $\sum_{k=1}^{I_c} \lambda_k^2$. The indicators "dry" or "wet" refer to whether the TLM (and therefore corresponding adjoint model) consider only dry or also moist

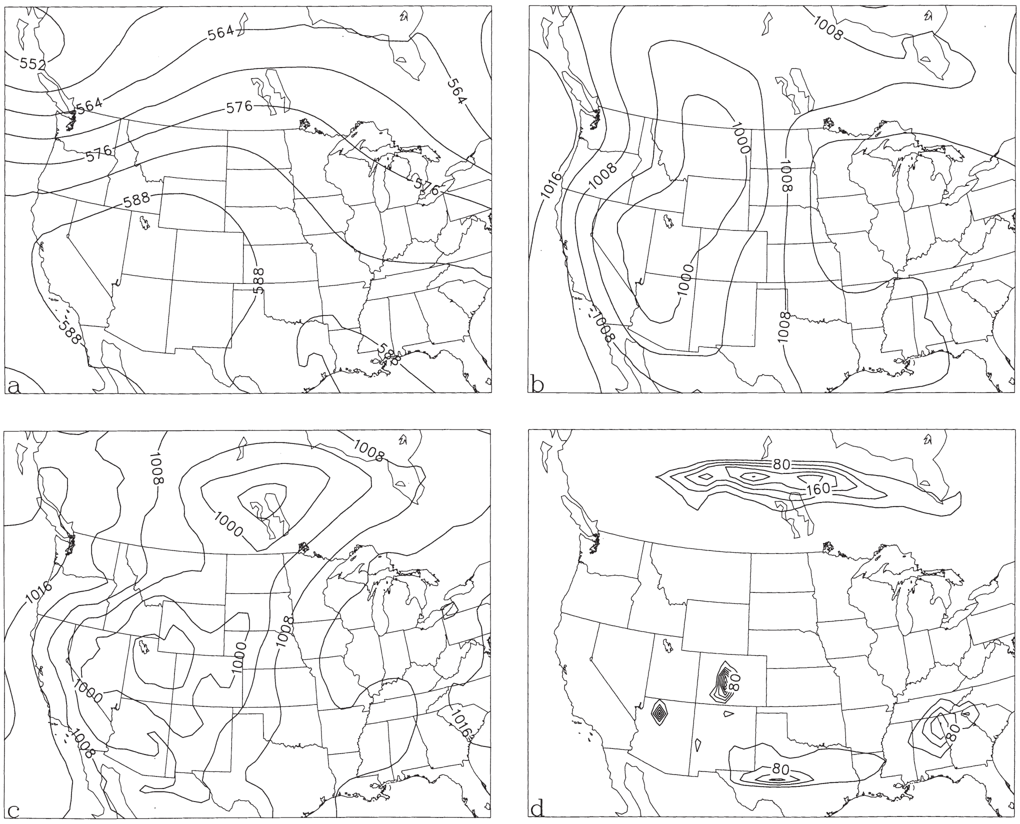

Fig. 2. As in Fig. 1, except for case S2.

Table 1. *Description of experiments and some results*

Case no.	Case label	Physics	Norm	N	I	I_c	N_g	λ_1^2	$\lambda_{I_c}^2$	Sum
1	W1	dry	E	79354	1000	406	—	54.9	1.53	1552.3
2	W1	wet	E	79354	600	257	—	322.0	2.23	2387.7
3	W1	dry	R	24940	600	265	270	38.1	1.01	791.6
4	W1	dry	TR	4830	390	180	169	28.2	0.95	504.1
5	W1	wet	TR	4830	750	416	175	39.5	0.35	762.2
6	W2	dry	TR	4830	400	191	171	19.6	0.90	521.8
7	W2	wet	TR	4830	400	197	178	59.2	0.90	633.4
8	S1	dry	E	65080	600	258	—	64.9	1.53	1117.3
9	S1	wet	E	65080	600	276	—	2547.9	1.55	6627.4
10	S1	dry	TR	3969	400	213	104	20.4	0.44	381.9
11	S1	wet	TR	3969	400	211	115	368.2	0.47	1126.3
12	S2	dry	TR	3969	400	196	124	11.1	0.62	354.7
13	S2	wet	TR	3969	400	206	134	81.1	0.61	659.2

N is the size of the (sub-) space measured by the norm. I is the number of iterations of the Lanczos algorithm performed. I_c is the number of converged λ^2. N_g is the number of $\lambda_k > 1$. Sum is $\sum_{k=1}^{I_c} \lambda_k^2$.

physical processes. N_g is the number of growing SVs determined with respect to the indicated norm and case. The sum over λ_k^2 is used in Section 5 to bound predictability estimates.

(a) *Case W1*

The converged λ_k^2 for the 5 sets of SVs obtained for case W1 appear in Fig. 3. The leading SVs for these sets are discussed in Ehrendorfer et al. (1999) and Errico (2000a, b). First, consider set 1 obtained for the E-norm with dry TLM (curve ED in Fig. 3). The λ_k^2 decrease by a factor of 10 between $k = 1$ and 52, but by only a factor of 3.5 between $k = 52$ and 400. The curve

$$\lambda_k \propto k^{-1/3} \tag{4.1}$$

also appears in Fig. 3, indicating that (4.1) is a fairly good approximation for this particular set, except for the leading few SVs. (No generality of approximation (4.1) to other cases, models, or norms is claimed.) Extrapolating using (4.1) suggests that the first decaying mode is near $k = 770$ which would mean that approximately 1% of the SVs are growing ones. At least 2000 Lanczos iterations may be required to confirm such a result.

Set 2 is obtained for the E-norm with moist TLM (curve EM in Fig. 3). The leading value is about 6× larger than for set 1, due to moist convective instabilities absent in the dry TLM (Errico and Raeder, 1999). As k increases, however, λ_k^2 drops even faster than for the dry TLM, appearing to become asymptotic to the curve for set 1, so that λ_{250}^2 for the two sets differ by only a factor of 1.1. It is therefore not apparent that the number of growing SVs is much larger than for the dry case.

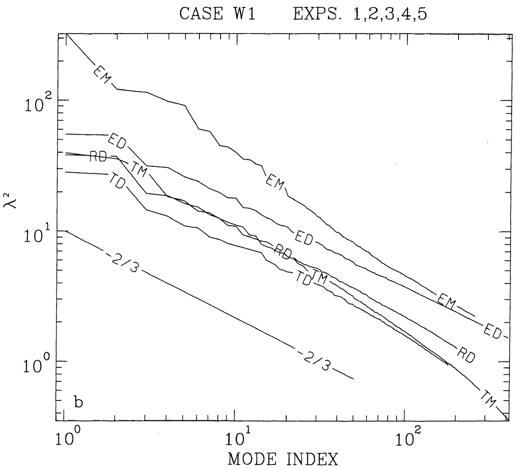

Fig. 3. Converged λ_k for the five sets of SVs determined for case W1. E, R, and T denote results for the energy, rotational mode, and truncated rotational mode norms, respectively. M and D denote results for moist and dry TLM, respectively. A line for which $\lambda_k^2 \propto k^{-2/3}$ is also indicated.

The effects of replacing the E-norm by the R-norm on the leading SVs for case W1 are described in Errico (2000a). With the dry TLM, the lack of consideration of gravitational modes by the R-norm, especially at the initial time, reduces λ_k^2 compared with the E-norm values for the same k (by as much as 30%). With the moist TLM, the reduction of λ_1^2 compared with the corresponding E-norm result is even greater as the convectively driven SVs determined for the latter set are excluded because they are dominantly gravitational (Errico, 2000b).

The effect of the R-norm beyond the leading SVs is also profound. For the non-truncated R-norm with the dry TLM, 270 SVs are growing, which is only 1.1% of total number (N) of SVs that span the subspace considered by the norm. When the R-norm is truncated, the number of growing modes is further reduced by 27%, but since N is also reduced by a factor of 5.2, 3.5% of the SVs considered by this norm grow. For the R- and TR-norms, the number of growing modes has been determined explicitly (no extrapolation of results required) since a sufficient number of SVs have been determined.

For large k, the curves for the R- and TR-norms are concave downward (on the log–log plot), reflecting that the rate of decay of λ_k^2 with increasing k is itself accelerating. For k smaller than the values where this acceleration is apparent, and excluding a leading few SVs, (4.1) seems to be a good approximation for both R- and TR-norms with the dry TLM.

For the truncated R-norm with the moist TLM, there are only 6 more growing SVs than for the corresponding dry TLM result. As in the E-norm comparisons, it is therefore not the case that the spectrum of λ_k^2 rises significantly for all the growing SVs when moist physics is considered. Note that $\lambda_{416}^2 = 0.35$, indicating that many SVs are significantly damped, rather than near neutral.

Several SVs for large indices k were examined in greater detail with the dry TLM. The E-norm SV for $k = 219$ ($\lambda^2 = 2.26$) is such an example. The T' field on $\sigma = 0.65$ at the initial time for this SV appears in Fig. 4a. This is the field that contributes most to the initial value of the norm for this SV. Unlike the leading SV, it is not very localized. This field is initially near 0 in the eastern-most portion of the domain because any perturbation there would be subsequently advected out of the domain, effectively damping the norm. The v' field on $\sigma = 0.35$ at the final time appears in Fig. 4b. It is near 0 in the western-most portion of the domain because no perturbation is advected in from the inflow boundary. Initially, 45% of the E-norm is contributed by gravitational modes, and 47% at the final time, indicating a significant lack of dynamic balance (as discussed in detail in Errico (2000b). This result confirms the speculation of Ehrendorfer and Errico (1995) that such unbalanced, slowly growing SVs may exist, greatly increasing the number of growing SVs for the E-norm. The higher-indexed R-norm SVs have the same character as the E-norm ones, in the sense that they are non-local, initially absent in the east and finally absent in the west, but of course they are geostrophic at initial time. Nothing apparently

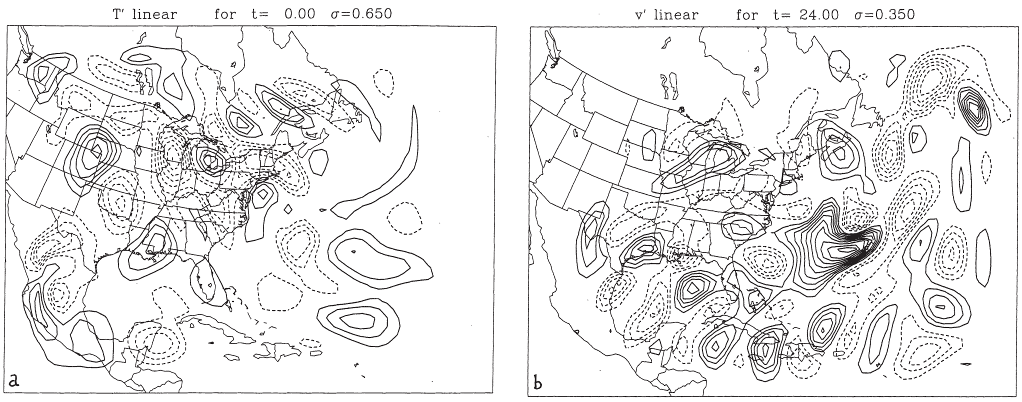

Fig. 4. The (a) initial T' field on $\sigma = 0.65$ and (b) final v' field on $\sigma = 0.35$ for SV number 219 determined for the E-norm, case W1, with the dry TLM. Contour intervals are respectively 0.25 K and 0.5 $\mathrm{m\,s^{-1}}$, with zero-contours omitted.

distinguishes the structures of the slightly growing SVs from those slightly damping.

(b) *Other cases*

Four sets of λ_k^2 spectra for case S1 appear in Fig. 5a. The leading SVs for the *E*-norm sets for case S1 are discussed in Ehrendorfer et al. (1999). Two sets of λ_k^2 for case W2 and two for case S2 appear in Fig. 5b. Only the TR norm was considered for cases W2 and S2.

There are many similarities between the results for case W1 compared with those for the other cases. In particular, when individual norms and cases are considered, leading SVs with the moist TLM are larger than with the dry TLM, but as k increases, the 2 λ_k appear to asymptote to each other in the ranges of values determined. Also, for all sets with the dry TLM, (4.1) is a good approximation within a broad range of k (perhaps less broad for case S2). For large k, the *R*-norm curves are concave downward on the log–log plot. In all cases, only a small percentage of *R*-norm modes are growing ones.

The numbers of growing *TR*-norm SVs are smaller in both summer cases compared to the winter ones. This may be partly due to N being smaller for the summer cases, although the % of growing SVs is also less. The numbers of growing SVs may also be reduced for the summer cases because their domains are smaller, but the wind speed of the reference states are also smaller, which should diminish the area in which perturbations are damped by the lateral boundaries, acting to mitigate this reduction.

Unlike for case W1, the *TR*-norm λ_k^2 in each of the other cases are much larger for the moist TLM compared to the dry one. For cases S1, S2, and W2, the ratio of moist to dry values are 18, 7.4, and 3, respectively, compared with 1.3 for case W1. Case W1 may therefore be an anomalous one. The large ratios for the summer cases are likely due to the summer continental air being both more convectively unstable and more moist.

For each case, the structure of the leading SV for the *E*-norm with the moist TLM is very different from that with the dry TLM, as discussed by Ehrendorfer et al. (1999) and Errico (2000b). In particular, the leading, moist TLM SV has smaller horizontal scale, has a different location, and is much more gravitational at both initial and final times compared with the corresponding result for the dry TLM. Similar differences apply to the contrasting *TR*-norm SVs, except with regard to increased initial gravitational energy since that energy is constrained to be zero in the initial *TR*-norm SVs. Although the *TR*-norm SVs also do not include measures of energy in gravitational modes at the final time when computing growth, they do generally create gravitational components as they evolve in the TLM.

The initial T' fields at $\sigma = 0.85$ and $\sigma = 0.65$,

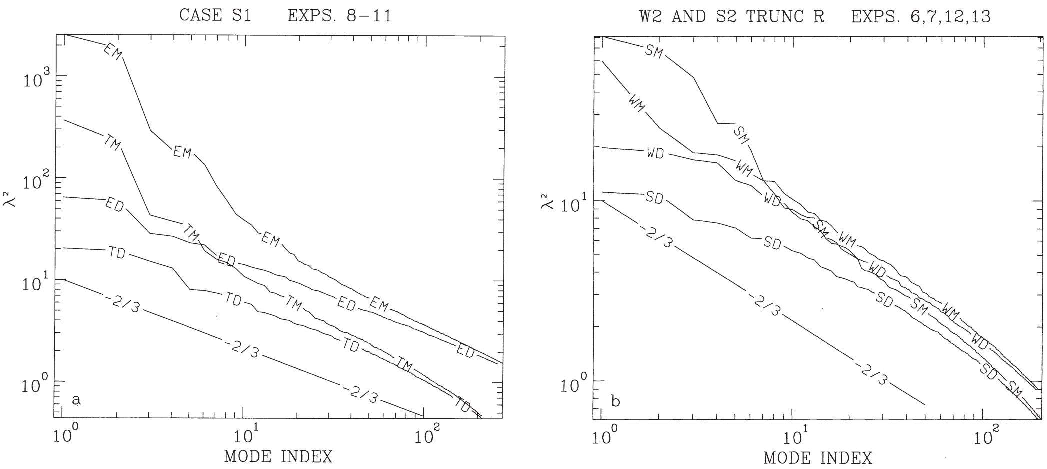

Fig. 5. As in Fig. 3, except for (a) case S1 and (b) cases S2 and W2. In (b), only the truncated rotational mode norm has been used, and the first letter designating the curves indicates the (S) S2 or (W) W2 case.

respectively, and the final v' fields at $\sigma = 0.25$ and $\sigma = 0.85$, respectively, for the S2 moist and dry TLM results for the leading SV for the TR-norm appear in Fig. 6. The chosen σ-levels are those where the perturbation variances are the largest, and the final time result includes the ageostrophic portion of the perturbation that evolves from the initial-time SV. The two SVs are dominant in different locations, and the horizontal scale for the moist results are smaller than for the dry ones. At both initial and final times, the perturbations peak on different σ-levels. At the final time, the ratio of rotational to gravitational contributions to E is 5 : 1 for the dry TLM results, but half that for the moist results. If the leading moist SV is evolved in the dry TLM, it grows by only a factor of 2.8 measured by the R-norm. So, for this case, the moist physics has a profound effect on the SVs, even without the introduction of dominantly ageostrophic structures, although not many additional growing structures are obtained.

A contrasting result is obtained for case S1 for the same norm. For it, the final-time contributions to E by rotational and gravitational components of the leading SV is nearly 1 : 1. Ageostrophic effects therefore appear to be important for the growth of this SV.

5. Predictability implications

Any initial perturbation to the TLM initial conditions can be described as:

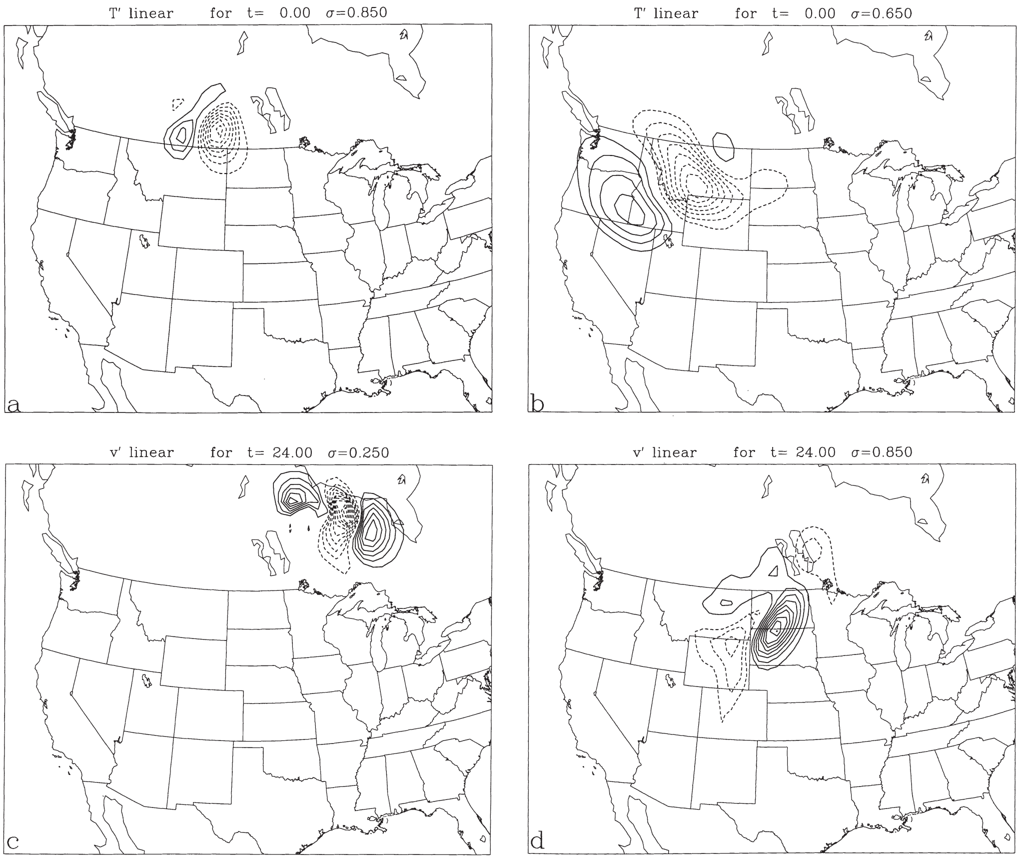

Fig. 6. The (a, b) initial T' fields and (c, d) final v' field on the indicated σ-levels for the leading SV determined for the TR-norm, case S2, with the (a, c) moist and (b, d) dry TLMs. Contour intervals are 0.2 K, 0.25 K, 5 ms^{-1}, and 2 ms^{-1} for (a–d), with zero-contours omitted.

$$\boldsymbol{x}(0)=\sum_{k=1}^{N}\alpha_k\boldsymbol{x}_k(0), \tag{5.1}$$

where $\boldsymbol{x}_k$ are the SV initial structures and α_k are their respective scalar coefficients. Since the SVs form an orthonormal basis with respect to the norm (2.1) at both initial and final times,

$$\|\boldsymbol{x}(0)\|^2=\sum_{k=1}^{N}\alpha_k^2, \tag{5.2}$$

$$\|\boldsymbol{x}(t)\|^2=\sum_{k=1}^{N}\alpha_k^2\lambda_k^2, \tag{5.3}$$

using the norm for which the SVs are determined.

Now consider that $\boldsymbol{x}(0)$ is selected randomly from a probability distribution that is white with respect to a norm specified by (2.1); i.e., each component of $\boldsymbol{z}=\boldsymbol{N}^{1/2}\boldsymbol{x}(0)$ is randomly selected from the same distribution, such that the variance of each component is identical, distinct components are uncorrelated, and the mean of each distribution is zero. Since the SVs form an orthogonal basis with respect to the norm defined by $\boldsymbol{N}$, the expected values of each α_k^2 are also all identical for this probability distribution. The ratio η of the expected values of $\|\boldsymbol{x}(t)\|^2$ to $\|\boldsymbol{x}(0)\|^2$ is therefore simply $N^{-1}\sum_k\lambda_k^2$. This η may also be interpreted as the expected growth $\|\boldsymbol{x}(t)\|^2/\|\boldsymbol{x}(0)\|^2$ of these random perturbations if an additional constraint is applied requiring that, for each random realization, the initial variance is normalized to some constant.

It is important to note that whether growth of a random initial perturbation should be expected depends not just on λ_1 or even the first few λ. If N is very large, then the probability that a random perturbation projects strongly on any single or small set of SVs is correspondingly small. Even though subsequent growth in the leading SVs will be greater than growth in other individual SVs, the overall behavior can be dominated by the much larger set of non-leading SVs. This is one of the reasons why knowledge of the entire spectrum of SVs should be of interest.

Values of η cannot be determined from our results because not all the values of λ_k are known. Upper and lower bounds can be established, however, by assuming that all unknown (i.e., non-converged) values are equal to the smallest known value or to 0, respectively.

Bounds on η for each experiment appear in Table 2. Those for the E-norm are not useful because so few converged SVs have been determined. Even the spread between upper and lower bounds for some of the R- and TR-norm experiments is broad, for the same reason. For the W2, S1, and S2 TR-norm cases, the upper bound is less than 1 and, for set 5, it is only 0.48. In these latter cases, the η therefore actually indicate that decay should be expected.

Ehrendorfer and Tribbia (1997) and Farrell (1990) demonstrate that if perturbations are interpreted as possible errors and the SVs are determined using a norm where N is equivalent to the inverse of an error covariance matrix, then the final-time SVs may be interpreted as eigenvectors of the final-time error covariance measured with respect to the selected norm. These are also known as the empirical orthogonal functions (EOFs) or principal components describing the covariance statistics of the final state. The λ_k^2 describe the variance contributed by each EOF. The ratio

$$\mu_m=\frac{\sum_{k=1}^{n}\lambda_k^2}{\sum_{k=1}^{N}\lambda_k^2} \tag{5.4}$$

is therefore the fraction of final-time variance explained by the first n EOFs. As for η, the values of μ_{n} depend on the entire set of λ_k and knowledge of that set should be of interest.

It is difficult to justify the use of the E-norm as a proxy for an analysis error covariance weighted norm. The TR-norm is arguably a more appropriate norm because at least it corresponds to spatially and dynamically correlated fields (Errico, 2000a). Predictability experiments have been conducted using random perturbations that are white with respect to covariances similar to those corresponding to E- and (independently) TR-norms (Errico and Baumhefner, 1987). So, although it is difficult to justify our E- and TR-norm μ_n as describing forecast error covariances (requiring an analysis error covariance norm at the initial time), our use of μ_n does describe ensemble error covariance fractions relevant to past predictability studies.

Since the sum in the denominator of (5.4) is unknown from our results, μ_n cannot be determined. An upper bound can be established, however, by summing over only the known portion of the spectrum. Lower bounds are established by

Table 2. *Bounds estimated for η and μ_{10} for each experiment*

Case no.	Case label	Physics	Norm	η upper	η lower	μ_{10} upper	μ_{10} lower
1	W1	dry	E	1.54	0.02	0.19	0.00
2	W1	wet	E	2.25	0.03	0.41	0.01
3	W1	dry	R	1.03	0.03	0.16	0.01
4	W1	dry	TR	1.01	0.10	0.27	0.03
5	W1	wet	TR	0.48	0.16	0.27	0.09
6	W2	dry	TR	0.97	0.11	0.26	0.03
7	W2	wet	TR	0.99	0.13	0.31	0.04
8	S1	dry	E	1.54	0.02	0.26	0.00
9	S1	wet	E	1.65	0.10	0.82	0.05
10	S1	dry	TR	0.51	0.10	0.29	0.05
11	S1	wet	TR	0.73	0.28	0.70	0.27
12	S2	dry	TR	0.67	0.09	0.21	0.03
13	S2	wet	TR	0.74	0.17	0.47	0.11

assuming all non-converged λ_k^2 are equal to the smallest converged value. Such bounds are presented in Table 2 for $n = 10$ for each set and in Fig. 7 for $n = 1, \ldots, 40$ for each set of cases S2 and W2.

Both upper and lower bounds for μ_i are typically larger for the moist sets compared to the corresponding dry ones, as should be inferred from the shapes of the spectra presented in Section 4. For the sets using the moist TLM, the upper bounds increase quickly with increasing n due to the very large λ_k for the first few k. For 3 of the 4 sets using the TR-norm with the moist TLM, the lower bounds indicate that at least 9% of the variance is explained by the first 10 leading SVs. For all the sets for the TR-norm with the dry TLM, $0.29 \geqslant \mu_{10} \geqslant 0.03$. The bounds do not provide a tight estimate because too few λ_i are known. All these bounds suggest, however, that even 0.25% of the possible SVs can explain a significant fraction of the forecast error variance.

6. Similarity measures

Although individual pairs of SVs produced using two distinct norms may appear dissimilar and uncorrelated, the subspaces spanned by a pair

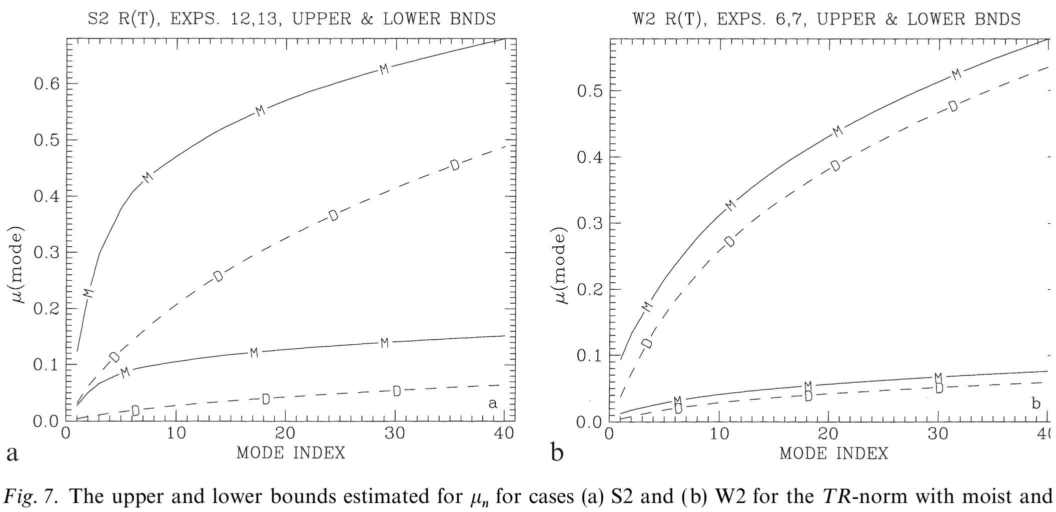

Fig. 7. The upper and lower bounds estimated for μ_n for cases (a) S2 and (b) W2 for the TR-norm with moist and dry TLMs (labeled as M and D, respectively).

of sets of several SVs for the two norms may be more similar. A measure of such similarity based on the projections of one set of SVs on another is described by Buizza et al. (1993). It is defined as

$$s_n(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{N}_s) = \frac{1}{n} \sum_{i,j=1}^{n} m_{i,j}(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{N}_s), \tag{6.1}$$

where $\boldsymbol{A}$ indicates a first set of SVs $(\boldsymbol{a}_1, \boldsymbol{a}_2, \ldots, \boldsymbol{a}_n)$, $\boldsymbol{B}$ indicates a second set $(\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n)$, $\boldsymbol{N}_s$ indicates the particular norm used to measure the similarity, and

$$m_{i,j} = (\boldsymbol{a}_i^{\mathrm{T}} \boldsymbol{N}_s \boldsymbol{b}_j)^2. \tag{6.2}$$

Here, $\boldsymbol{N}_s$ will always be the E-norm, and only SV structures at the initial time will be examined. Parallel sub-spaces have $s = 1$, and orthogonal sub-spaces have $s = 0$. Note that, although the sum in (6.1) cannot decrease with increasing n, s_n can. For further interpretation of this norm see Buizza (1994).

For reference, values of s_n were computed for pairs of experiments that were expected to have little commonality. These are pairs determined for the TR-norm on identical model grids, but for distinct synoptic cases. For all such pairs for the winter cases, s_n varies approximately linearly between 0 for $n = 1$ and 0.11 for $n = 100$. For the summer cases, the similarity between such pairs is typically larger for any value of n compared to the winter cases, but $s_n < 0.23$ for all n. Values of s_n smaller than for these pairs therefore indicate insignificant similarity.

The similarity index produced for the n-leading SVs for pairs of experiments for case W1 appear in Fig. 8a. Each curve is labeled with the experiment numbers compared. Curve 1–2 is the s_n comparing SVs produced using the E-norm for dry and moist linearizations. It is zero through $n = 3$ because the SVs produced by each model version are in different locations (as described in Ehrendorfer et al., 1999). As n increases, s_n typically increases, but not monotonically. $s_{19} = 0.5$, and $s_{100} = 0.707$, indicating significant and increasing similarity between the subspaces for $n \geqslant 19$.

The remaining curves in Fig. 8a compare SVs produced using different norms. Curve 1–3 compares dry results with the E- and R-norms. When only the 2 leading SVs are considered, $s = 0.74$. For increasing n there is a general, although not monotonic, decrease in s_n, such that $s_{100} = 0.69$. Curve 1–4 compares results for dry E- and TR-norms. It is similar to that for 1–3, except corresponding s_n are approximately 35% smaller. These results are explainable by the smaller subspaces spanned by the successive norms considered. Curve 2–5 also compares E- to TR-norm results, but for moist physics. It indicates less similarity than for the dry SVs, consistent with an hypothesis that ageostrophic or small scale fields not considered by the TR-norm may grow with respect to the E-norm, particularly when moist physics is considered.

Similarity indices for some other some pairs of experiments appear in Fig. 8b. The curves labeled 6–7, 8–9, 10–11, 12–13 each compare pairs of SVs produced using identical norms but dry versus moist linearizations. For experiments 6 and 7 (both using the TR-norm), $s_3 = 0.7$, and the similarity remains high for all n examined, with $s_{100} = 0.89$. Of the 4 comparisons in Fig. 8b, the comparison indicating the least similarity is for experiments 8, 9 using the E-norm, but even for this, $s_n > 0.5$ for $n > 27$ and $s_{100} = 0.7$. In contrast, the s_n for experiments 8, 10 and 9, 11 (not shown) remain below 0.38 for $n \leqslant 100$ and the curves are rather flat. Results for these cases therefore are qualitatively similar to those for case W1 and the same general characterization applies.

Further examination of individual values of $m_{i,j}$ reveals that higher order SVs produced using the E-norm have little in common with individual SVs produced using the TR-norm for the same case and physics (the maximum value of $m_{i,j}$ for each $j > 50$ considering all i is typically less than 0.08. In contrast, when $m_{i,j}$ are examined for SV pairs produced for the TR-norm applied to identical cases, but with different physics, even for $j > 50$, the maximum $m_{i,j}$ is often greater than 0.25. This too is consistent with the characterization of higher-indexed E-norm SVs as having large ageostrophic components. Also, the large similarity between higher-indexed TR-norm SVs produced with different physics suggests that the moist physics for these cases can be characterized as having essentially a modulating effect on SVs if the initially ageostrophic, convectively driven SVs are excluded, as discussed by Errico and Raeder (1999) for some leading SVs.

7. Conclusions

For the 13 sets of SVs described, encompassing 4 synoptic cases, 3 norms, and 2 versions of the

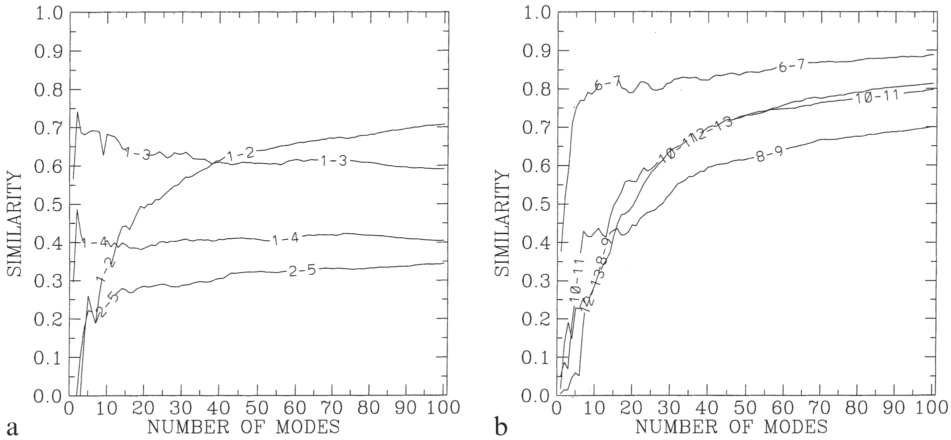

Fig. 8. The similarity index as a function of number of SVs considered for the indicated pairs of experiments: (a) for case W1, (b) for other cases.

TLM, only small fractions of the possible number of SVs that span the spaces measured by the norms appear to be growing. Although the actual fraction has not been determined for the energy norm results, extrapolation suggests that approximately 1% are growing with respect to that norm. For the rotational mode norm, the same percentage was obtained (without the need for extrapolation), although the number of SVs spanning the subspace measured by that norm is less than 1/3 that of the *E*-norm. The many additional growing SVs obtained with the *E*-norm are very slowly growing structures, some with large gravitational mode components. When perturbations are intended to represent possible analysis errors and those analysis errors are considered to be dominantly geostrophic, these structures should be of little interest, and the *R*-norm can be used to appropriately filter them.

When the measuring of small vertical and horizontal scales of fields is excluded by using the truncated *R*-norm, up to 4% of the SVs are found to grow. The increase of percentage compared with the *R*-norm without truncation is due to the factor of 5 decrease in the number of SVs possible, since most of those small scale structures excluded due to the truncation are expected to be primarily damping ones. The truncation also excludes some growth and reduces the maximum growth. If such effects are undesirable, this norm should not be used, unless it could be justified as being a reasonable approximation to an inverse, analysis error covariance weighted norm. It is also easy to make the truncation less severe, particularly with regard to vertical scales.

When the dry TLM is used, the singular values are approximately proportional to $k^{-1/3}$ over a wide range of mode index k within the range of growing SVs. We do not know how general this relation is, however, since even in the cases examined here, the fit to it varies considerably. Although the leading values are larger with the moist versus dry TLM, the moist values appear to asymptote toward the corresponding dry ones as k increases within the range of values investigated. The numbers of growing SVs are therefore not much different for the moist and dry physics. The primary effect of the moist physics on the spectrum of λ_k^2 is to increase the leading singular values, primarily by adding a few new fast-growing structures (Ehrendorfer et al., 1999). For SVs determined for the *E*-norm, this increase is huge, but even for the truncated rotational mode norm, it still can be quite large: an increase by a factor of 4.3 was obtained for SV1 for a summer case. At the final time, that SV had a large gravitational component (which is not, however, measured by the *R*-norm).

The spaces spanned by the leading SVs produced for identical cases using identical norm, but different (moist versus dry) physics are very similar except for the first several modes. When spaces of the leading 100 SVs are compared, the similarity index is between 0.7 and 0.9. In contrast, when such spaces are compared for SVs produced for identical cases and physics, but using E- versus TR-norms, they are not very similar (index less than 0.41). For such pairs of experiments when moist physics is considered, the similarity is even less, consistent with the hypothesis that perturbations in ageostrophic and small scale fields not considered by the TR-norm can lead to growth of the E-norm when moist physics is considered.

Since the number of decaying SVs so dominates the number of growing ones, the expected result of a random, white-noise perturbation within the subspace considered by a norm, when normalized at the initial time so that its norm is a prescribed constant, is that the perturbation will subsequently damp. Since the singular values decay so rapidly with increasing k, however, the percentage of final-time variance for an ensemble of such perturbations can be explained by even a small number of leading SVs. For example, for the TR-norm, 10 SVs explain more than 27% of the variance of such perturbations in one of the summer, moist TLM cases. For this reason, SVs can be very efficient for characterizing likely forecast errors due to initial condition errors that initially have a white-noise distribution with respect to the norm as demonstrated by Ehrendorfer and Tribbia (1997) for a much simpler model. For this to be true, however, it is also necessary that the perturbations behave sufficiently linearly.

Since the lateral boundaries are not perturbed, their "sweeping effect" described by Errico and Baumhefner (1987) acts to both restrict and damp some perturbations. This sweeping can be seen in the near neutral SVs. If significantly growing structures move along with the synoptic structures of the reference state (about which the linearization is performed), then this effect should be limited to less than 1/3 the domain (considering a propagation rate of $15\ \mathrm{m\,s^{-1}}$). This means, however, that perhaps no more than 67% of the possible SVs can grow.

The linearization of the MAMS convective scheme is capable of producing some extraordinarily large perturbation growth, as discussed in Errico and Raeder (1999). Such growth is not observed when perturbations are introduced in the nonlinear model with initial amplitudes the size of analysis uncertainty. The scheme's linearization has therefore been modified to remove TLM solutions for which the norm can more than double within 1 h, so that the linearized results are more comparable with the corresponding nonlinear ones. This has the effect of removing some potentially fast growing SVs. As long as our primary interest is in the underlying nonlinear problem, however, the results presented here are arguably a better indicator of growing, relevant SVs. Also, it should be noted that even with the convective filter, very large growth is still obtained (e.g., growth by a factor of 2547 in case S1).

Fillion and Mahfouf (2000) have shown that linearizations of different parameterized moist convection schemes can produce very different results. It is therefore unknown to what extent many specific differences between moist and dry results reported here may be generalizable to other models. Also, our results may depend strongly on the nature and tuning of our model's dissipation mechanisms. Without much additional experience or confirmation of our results by others, with their own models, indeed there is little we can say about the generality of our results. We have been able to show, however, some consistency between cases we investigated and with other diagnostics or experiments (e.g., with Errico (2000b), or Errico and Baumhefner (1987), as discussed further by Ehrendorfer and Errico (1995)).

None of the singular value spectra have the symmetries with respect to k obtained for much more weakly damped, low-resolution quasi-geostrophic models (Reynolds and Palmer, 1998). In particular, there are many times more damping than growing SVs. The modeled system, as typical of other models suitable for numerical weather prediction, is therefore not nearly a Hamiltonian system, and therefore the properties of such systems discussed by Pu et al. (1997) do not apply.

The truncated rotational mode norm is shown to have some desirable properties. For the leading SVs, some of these properties have been discussed in Errico (2000a). Additionally, here it has been shown that much larger growth can be obtained with moist versus dry physics, although the norm excludes small-scale gravitational perturbations at the initial time. Compared with the E-norm, the truncated rotational mode norm excludes several hundred, near neutral, mixed rotational-gravitational SVs, and reduces the size of the subspace considered for the SVs by a factor of 15. For this reason, a large percentage of the SVs may be obtained with a reasonable number of Lanczos iterations and even some damping SVs may be readily found.

7. Acknowledgments

The authors thank C. Reynolds for reviewing early versions of the manuscript. This work was partially funded by US Office of Naval Research contracts N00173-96-MP-0086 and N00014-99-1-0017 and NASA contract W-18,077, Mod. 6. M. Ehrendorfer was supported by the Austrian FWF (Fonds zur Förderung der wissenschaftlichen Forschung) project P13729-GEO.

REFERENCES

Barkmeijer, J., Buizza, R., Palmer, T. N., Puri, K. and Mahfouf, J.-F. 2001. Tropical singular vectors computed with linearized diabatic processes. *Quart. J. Roy. Meteor. Soc.* **127A**, in press.

Bourke, W. and McGregor, J. L. 1983. A nonlinear vertical mode initialization scheme for a limited area prediction model. *Mon. Wea. Rev.* **111**, 2285–2297.

Buizza, R. 1994. Sensitivity of optimal unstable structures. *Quart. J. Roy. Meteor. Soc.* **120**, 429–451.

Buizza, R., Tribbia, J., Molteni, F. and Palmer, T. N. 1993. Computation of optimal unstable structures for a numerical weather prediction model. *Tellus* **45A**, 388–407.

Cohn, S. E. and Todling, R. 1996. Approximate data assimilation schemes for stable and unstable dynamics. *J. Meteor. Soc. Japan* **96**, 63–75.

Davies, H. C. and Turner, R. E. 1977. Updating prediction models by dynamical relaxation: an examination of the technique. *Quart. J. Roy. Meteor. Soc.* **103**, 225–245.

Ehrendorfer, M. 1997. Predicting the uncertainty of numerical weather forecasts. *Meteorol. Zeitschrift* **6**, 147–183.

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

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 error covariances through singular vectors. *J. Atmos. Sci.* **54**, 286–313.

Errico, R. M. 1997. What is an adjoint model. *Bull. Am. Meteor. Soc.* **78**, 2577–2591.

Errico, R. M. 2000a. Interpretations of the total energy and rotational energy norms applied to determination of singular vectors. *Quart. J. Roy. Meteor. Soc.* **126A**, 1581–1599.

Errico, R. M. 2000b. The dynamical balance of singular vectors in a primitive equation model. *Quart. J. Roy. Meteor. Soc.* **126A**, 1601–1618.

Errico, R. M. and Baumhefner, D. P. 1987. Predictability experiments using a high-resolution limited-area model. *Mon. Wea. Rev.* **115**, 488–504.

Errico, R. M. and Raeder, K. D. 1999. An examination of the accuracy of the linearization of a mesoscale model with moist physics. *Quart. J. Roy. Meteor. Soc.* **125**, 169–195.

Errico, R. M., Raeder, K. and Vukićević, T. 1994. Description of the NCAR mesoscale adjoint modeling system version 1 (MAMS1). NCAR Technical Note, NCAR/TN–410+IA. (Available from the National Center for Atmospheric Research, PO Box 3000, Boulder, CO 80307, USA.)

Farrell, B. F. 1989. Optimal excitation of baroclinic waves. *J. Atmos. Sci.* **46**, 1193–1206.

Farrell, B. F. 1990. Small error dynamics and the predictability of atmospheric flows. *J. Atmos. Sci.* **47**, 2409–2416.

Farrell, B. F. and Ioannou, P. J. 1996a. Generalized stability theory. Part I: Autonomous operators. *J. Atmos. Sci.* **53**, 2025–2040.

Farrell, B. F. and Ioannou, P. J. 1996b. Generalized stability theory. Part II: Nonautonomous operators. *J. Atmos. Sci.* **53**, 2041–2053.

Fillion, L. and Mahfouf, J.-F. 2000. Coupling of moist convective and stratiform precipitation processes for variational data assimilation. *Mon. Wea. Rev.* **128**, 109–124.

Gelaro, R., Buizza, R., Palmer, T. N. and Klinker, E. 1997. Sensitivity analysis of forecast errors and the construction of optimal perturbations using singular vectors. *J. Atmos. Sci.* **55**, 1012–1037.

Gelaro, R., Langland, R. H., Rohaly, G. D. and Rosmond, T. E. 1999. An assessment of the singular vector approach to targeted observing using the FASTEX data set. *Quart. J. Roy. Meteor. Soc.* **125**, 3299–3327.

Grimes, R. G., Lewis, J. G. and Simon, H. D. 1994. A shifted block Lanczos algorithm for solving sparse symmetric generalized eigenproblems. *SIAM J. Matrix Analysis and Applications* **15**, 228–272.

Mahfouf, J.-F. 1999. Influence of physical processes on the tangent–linear approximation. *Tellus* **51A**, 147–166.

Moorthi, S. and Suarez, M. J. 1992. Relaxed Arakawa–Schubert: A parameterization of moist convection for general circulation models. *Mon. Wea. Rev.* **120**, 978–1002.

Mullen, S. L. and Baumhefner, D. P. 1989. The impact of initial condition uncertainty on numerical simulations of large-scale explosive cyclogenesis. *Mon. Wea. Rev.* **117**, 2800–2821.

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., Kalnay, E., Sela, J. and Szunyogh, I. 1997. Sensitivity of forecast errors to initial conditions with a quasi-inverse linear method. *J. Atmos. Sci.* **125**, 2479–2503.

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

Strang, G. 1986. *Introduction to applied mathematics.* Wellesley-Cambridge Press, 758 pp.

Talagrand, O. 1981. A study of the dynamics of four-dimensional data assimilation. *Tellus* **33**, 43–60.