



Aggressive shadowing of a low-dimensional model of atmospheric dynamics

Ross M. Lieb-Lappen^{a,*}, Christopher M. Danforth^{a,b,c}

^a Department of Mathematics and Statistics, University of Vermont, Burlington, VT, United States

^b Vermont Advanced Computing Core, University of Vermont, Burlington, VT, United States

^c Vermont Complex Systems Center, University of Vermont, Burlington, VT, United States

ARTICLE INFO

Article history:

Received 6 May 2011

Received in revised form

15 November 2011

Accepted 1 December 2011

Available online 9 December 2011

Communicated by T. Sauer

Keywords:

Shadowing

Stalking

Ensemble forecasting

Inflation

ABSTRACT

Predictions of the future state of the Earth's atmosphere suffer from the consequences of chaos: numerical weather forecast models quickly diverge from observations as uncertainty in the initial state is amplified by nonlinearity. One measure of the utility of a forecast is its *shadowing time*, informally given by the period of time for which the forecast is a reasonable description of reality. The present work uses the Lorenz '96 coupled system, a simplified nonlinear model of atmospheric dynamics, to extend a recently developed technique for lengthening the shadowing time of a dynamical system. Ensemble forecasting is used to make forecasts with and without *inflation*, a method whereby the ensemble is regularly expanded artificially along dimensions whose uncertainty is contracting. The first goal of this work is to compare model forecasts, with and without inflation, to a true trajectory created by integrating a modified version of the same model. The second goal is to establish whether inflation can increase the maximum shadowing time for a single optimal member of the ensemble. In the second experiment the true trajectory is known a priori, and only the closest ensemble members are retained at each time step, a technique known as *stalking*. Finally, a *targeted inflation* is introduced to both techniques to reduce the number of instances in which inflation occurs in directions likely to be incommensurate with the true trajectory. Results varied for inflation, with success dependent upon the experimental design parameters (e.g. size of state space, inflation amount). However, a more targeted inflation successfully reduced the number of forecast degradations without significantly reducing the number of forecast improvements. Utilized appropriately, inflation has the potential to improve predictions of the future state of atmospheric phenomena, as well as other physical systems.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Society is often dependent upon the ability of scientists to accurately forecast the future state of chaotic physical systems, such as the weather. In extreme cases, atmospheric scientists are asked to anticipate natural disasters with ample time for the public to adequately prepare. To predict the likelihood of atmospheric phenomena, meteorologists estimate Earth's future atmospheric state using sophisticated computer simulations. In general, these mathematical models struggle to predict the specific behavior of chaotic physical systems due to (1) uncertainty in the initial state, (2) chaos, and (3) model error [1].

The initial state is an estimate of the condition of the atmosphere at the beginning of a forecast. Although weather monitoring devices observe much of the planet, for many areas (e.g. southern hemisphere oceans, and the upper atmosphere) the

measurements are spatially sparse and typically three orders of magnitude less in number than the degrees of freedom in a global weather model. Therefore, meteorologists must combine observations with prior forecasts to estimate initial conditions in a process called *data assimilation*.

Since perfect knowledge of atmospheric tendencies is unachievable, a great deal of recent research has focused on data assimilation, the process by which observations are combined with model predictions to give the best possible initial state, or *analysis* [2–6]. It is the analysis that is typically used as a proxy for the true state of the atmosphere at any time in the past. Although the analysis has inherent uncertainty from lack of perfect observations and an imperfect model, the modeler's goal is to create a forecast from this given state that remains reasonable (i.e. close to observations) for the longest duration of time. In the context of weather predictions, forecasters strive to be accurate for two weeks, the limit imposed by chaos. Today, five-day forecasts are as good as three-day forecasts were 30 years ago [7].

Despite these efforts, the mischaracterization of Earth's atmospheric state introduces uncertainty. In addition, the atmosphere

* Corresponding author.

E-mail addresses: rlieblappen@gmail.com (R.M. Lieb-Lappen), chris.danforth@uvm.edu (C.M. Danforth).

is a chaotic system, causing these small uncertainties in the initial state to grow exponentially in time. Finally, since the model itself is not a perfect representation of reality, inaccurate changes forecast by the model, namely model error, compound the uncertainty and lead forecasts to diverge quickly from observations. As Lorenz noted in 1965, the limit of predictability of the atmosphere is about two weeks, even with nearly perfect knowledge of the current state [8].

To account for initial state uncertainty, an accepted technique is to use *ensemble forecasting*, where a collection of perturbations to the best guess are chosen randomly from a distribution that reflects the measurement uncertainty and system dynamics near the given initial state [9]. Each ensemble member is forecast forward in time, yielding a collection of final states. This collection roughly represents the probability distribution of the model's forecast with associated uncertainty. For example, if 60% of the ensemble members predict rain, the forecaster assigns a 60% chance of rain. Given the limitations, the modeler's goal can be stated as trying to keep some ensemble members close to the observed truth for as long as possible, analogous to finding a shadowing trajectory for Earth's atmosphere.

Traditional mathematical shadowing theory was developed for hyperbolic systems, based on the Shadowing Lemma of Anosov [10] and formalized by Bowen [11]. Given a pseudo-trajectory of a model (i.e. one very close to an actual one-step trajectory), this lemma establishes the existence of a true trajectory that remains close for an arbitrary period of time. Later research has extended the lemma for a wide variety of hyperbolic systems (e.g. [12–14]). For these systems, the number of expanding directions remains constant, i.e. the number of Lyapunov exponents greater than zero is constant throughout the state-space, and small perturbations in stable directions decay exponentially in time. However, most physical systems (e.g. Earth's atmosphere) are non-hyperbolic, i.e. the number of positive Lyapunov exponents fluctuates along a trajectory through state space. For these systems, there does not exist a trajectory that shadows the truth for an arbitrarily long time [15–17]. More recent work has been focused on finding shadowing trajectories for low-dimensional non-hyperbolic systems (e.g. general ordinary differential equations [18,19], the driven pendulum [20], the Hénon map [21]).

For the purposes of the present research, the time period for which a particular forecast is an accurate representation of reality is referred to as the *shadowing time*. Danforth and Yorke [22] proposed an aggressive approach to shadowing called *stalking* to increase the shadowing time for a given forecast. For a system with n -degrees of freedom, an n -dimensional sphere was used to encompass the initial ensemble members. Throughout the length of the forecast, these members experience expansion away from each other in some dimensions and contraction toward each other in others, depending on the local finite-time Lyapunov exponent in each dimension. The ensemble forecast can be approximated by an n -dimensional ellipsoid for a short time.

The idea of stalking is to artificially impose some additional uncertainty in the contracting directions (i.e. those with a negative local finite-time Lyapunov exponent) at periodic intervals throughout the forecast. This is accomplished by inflating the ellipsoid along axes parallel to the contracting dimensions. Thus, if the true state of the system happens to suddenly expand along a previously contracting direction, as happens in systems exhibiting *unstable dimension variability* [23], some ensemble members will remain relatively close to the true state. Stalking is not currently used in weather forecasting, but an alternative version called *variance inflation* is used in ensemble-based data assimilation to ensure the state estimation algorithm does not put too much faith in the model forecasts and ignore observations when creating the analysis [24].

For the purpose of this paper, the terms “shadowing” and “stalking” are reserved for experiments where the truth is known a priori, and the goal is to identify an optimal ensemble member (Section 5). Using the technique described above to create forecasts, without prior knowledge of the truth, is here denoted forecasting with and without inflation (Section 4). In Section 2, we present the non-hyperbolic '96 Lorenz system [25] and the model, and examine the effect of changing basic parameters. In Section 3, the process of creating the initial ensemble and making forecasts is explained. In Section 4, the inflation technique is introduced, and forecasting results (with and without inflation) are presented. In Section 5, we present the results from stalking experiments. In Section 6, we use targeted inflation to isolate the instances in which inflation is most likely to improve a forecast. Finally, the findings are discussed with respect to operational weather prediction in Section 7.

2. The Lorenz 1996 system

In numerical weather prediction (NWP), the true evolution of the Earth's atmosphere will ideally be a plausible member of an ensemble forecast [9]. In our context, ensemble forecasting is used to model the trajectory \mathbf{z}^a , the “truth”, of some meteorological quantity (e.g. temperature, pressure). In the present study, both the truth and the forecast were created using versions of a simplified nonlinear model given by Lorenz (1996) to represent the atmospheric behavior at I equally spaced locations on a given latitude circle [25]. This system has been used to represent weather related dynamics in many previous studies (e.g. [26–28]). The governing first-order differential equations are given by [25]:

$$\frac{dx_i}{dt} = x_{i-1}(x_{i+1} - x_{i-2}) - x_i + F - \frac{hc}{b} \sum_{j=(i-1)+1}^{ij} y_j, \quad (1a)$$

$$\frac{dy_j}{dt} = -cby_{j+1}(y_{j+2} - y_{j-1}) - cy_j + \frac{hc}{b} x_{\lfloor \text{floor}[(j-1)/I] + 1} \quad (1b)$$

for $i = 1, 2, \dots, I, j = 1, 2, \dots, II$, and $n = (J + 1)I$. Although n represents the full state space size, we are primarily interested in the values of the slow variables x_i . Thus, for the purposes of this study we consider I to be the dimensionality of the system.

The values x_i represent slowly changing meteorological quantities whose dynamics are described by Eq. (1a). Since the set of x_i corresponds to locations along a single latitude circle, the subscripts i and j are defined to be in a cyclic chain. That is, we define $x_{-1} = x_{I-1}$, $x_0 = x_I$, and $x_1 = x_{I+1}$, and similarly for j . Each x_i is then coupled to J quickly changing, small amplitude variables whose behavior is governed by Eq. (1b). For our experiments, we set $I = 4, 5, 6$ and $J = 16$ yielding $n = 68, 85$, and 102 variables. A schematic is shown in Fig. 1 where eight slow variables (x_i) are each coupled to four fast variables (y_j), reflecting a total of $n = 40$ degrees of freedom. Note that the dynamics of each x_i are dictated by neighboring x variables and the corresponding set of coupled y_j variables.

As a low-dimensional model of atmospheric dynamics, this system is ideal for several reasons; primary among them being the ability to tune the relative levels of nonlinearity, coupling of timescales, and spatial degrees of freedom. The nonlinear terms in Eq. (1a) are meant to represent advection, and conserve the total energy of the system. The linear term signifies a loss of energy either through mechanical or thermal dissipation. External forcing F is then added to prevent the total energy from decaying completely. For all experiments we set $F = 14$. Consistent with the literature, we set $c = 10$ and $b = 10$, which forces the fast variables to oscillate 10 times quicker than the slow variables [26,28,17]. The coupling parameter is set to a value of $h = 1$ for the

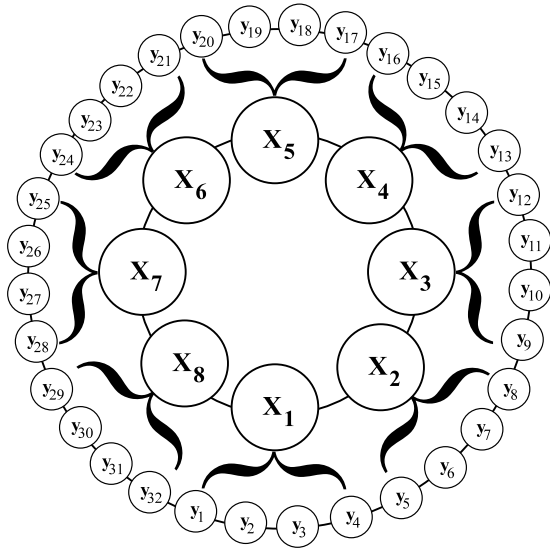


Fig. 1. Model schematic [26]. The $I = 8$ slow variables x_i can be thought of as locations along a latitude circle. Each variable has $J = 4$ corresponding y_j fast variables for a total of $n = 40$ degrees of freedom.

system truth, and $h = 0.5$ for the model. Note that one time unit corresponds to five days, the dissipative decay time [29].

Integration of the differential equations is completed using the fourth order Runge–Kutta method with a time step of 0.01. Rigorous shadowing attempts could be made using far more advanced methods of integration, with much smaller time steps [30]. However, for the purpose of this study of short forecasts, the difference is negligible.

2.1. Dynamics

The system dynamics can be observed by integrating equations (1a) and (1b) with $h = 1$, hereafter referred to as the *system*. In Fig. 2, a time series for a particular longitudinal profile ($I = 40, J = 16$) is shown after variable x_{13} is initially perturbed by five units. Profiles are recorded at 12-h intervals for days 0–5 (panel a) and days 50–55 (panel b). Advection is apparent as the energy from the perturbation is observed halfway around the latitudinal circle by day five. Lorenz and Emanuel [29] calculated a perturbation growth rate (doubling time) of approximately two days for this model, which agrees with trends in larger atmospheric models where errors double in roughly two days [31]. However, growth

rates over limited time intervals as in Fig. 2(b) can differ greatly depending on I and J .

Adjusting the parameters of the system, and examining the time series at a single latitudinal site, the interdependence of the slow and fast variable coupling is observed. For example, with $I = 6$ and $J = 16$, a “regular” oscillatory pattern occurs as an energy equilibrium is achieved between external forcing and dissipation. In this work, we will use the term “regular” to refer to these quasi-periodic patterns. However, as the significance of the fast variables is reduced by lowering the coupling parameter h , the time series becomes more complex. Less energy is able to dissipate from slow variables to fast, leading to greater irregularity in x_i .

Similarly, the regularity of a time series can be adjusted by varying the number of slow or fast variables. As illustrated in Fig. 3 (across columns), increasing J creates a more regular system for all values of I . For $I = 6$, the system is fairly regular for both $J = 16$ and $J = 40$. However, with only a few fast variables active (e.g. $J = 8$), oscillations become increasingly aperiodic. For $I = 8$, additional fast variables are required to achieve the consistent pattern; they allow the system’s energy to be more evenly distributed, and nonlinearities in the fast variables have less impact on the stability of the slow variables.

Adjusting the number of slow variables greatly alters the qualitative dynamics of the time series for each particular x_i (Fig. 3, across rows). For $J = 8$, only $I = 4$ exhibits regularity. However, with additional fast variables ($J = 16$), the $I = 6$ system exhibits the greatest regularity. For $I = 4$, the slow variable relative maxima are more pronounced than the relative minima, and the system appears to be quasi-periodic with a dominant frequency of roughly ten days. For $I = 6$, the symmetrical pattern emerges with a frequency of roughly five days. This pattern is similar when J is increased to 40, but with decreased amplitude. However, for $I = 4$ and $I = 8$ with $J = 40$, the pattern being repeated is more complex.

We make the preceding observations to motivate the choice of Lorenz ‘96 as an appropriate testbed with which to perform ensemble forecasting experiments. The dynamics are quite sensitive to the system parameters, and we will describe forecast results, which vary depending on the underlying dynamics.

2.2. Coupling

Each x_i in the system is coupled to J small-amplitude fast variables y_j , and the strength of this coupling is controlled by the parameter h . As mentioned above, the system is defined by setting $h = 1$. However, the model is rendered imperfect by dampening the effect of the fast variables in Eq. (1b) by setting $h = 0.5$,




Fig. 2. A time series of the system shows the effect of a five-unit perturbation. Panel (a) shows days 0–5 and panel (b) shows days 50–55. Perturbations appear to move primarily to the east (right). After five days, perturbation effects can be observed at over half of all the locations ($I = 40, J = 16$).




Fig. 3. A 50 day time series for x_3 illustrates the dependence of the system stability on the number of slow (I) and fast (J) variables. The number of slow variables varies from $I = 4$ (a–c) to $I = 6$ (d–f) to $I = 8$ (g–i). The fast variables vary from $J = 8$ (left column) to $J = 16$ (middle column) to $J = 40$ (right column). Generally, increasing J leads to increased regularity, while changes in I have an impact that is dependent upon the magnitude of J .





Fig. 4. A 20 day time series for x_1 and $y_{1,2,3,4}$ using $I = 8, J = 4$. For the top frame, the solid line represents the time series for x_1 with fast variable coupling set at $h = 1$ (the system). The dashed line has fast variable coupling reduced to $h = 0.5$ (the model). The bottom four frames are the fast variable time series for the system $y_{1,2,3,4}$ coupled to x_1 , as illustrated in Fig. 1.


hereafter referred to as the *model*. A 20 day time series for a single x_i with its corresponding y_j 's is shown in Fig. 4 ($I = 8, J = 4$). In the top frame we compare the time series of x_1 for integrations of the same initial condition by both the system and the model. By day 20 the x_1 values differ by the climatological span of x_1 . In addition to increasing the amplitude of slow variable oscillations, reducing the coupling h has the effect of creating a less regular time series. Note that the fast variables (shown for the system) vary with amplitudes on the order of 10% of the slow variables.

3. Methods


From a given initial condition, the trajectory of the truth (\mathbf{z}^a) is created by integrating the system in Eqs. (1a) and (1b). Two- and three-dimensional views of orbit segments of this attractor with $I = 4, J = 16$ are shown in Fig. 5(a) and (b). As mentioned above, the forecast for each ensemble member is then created by reducing the weight of the fast variables by 50% in the model, relative to the system. Two- and three-dimensional slices for the model can be seen in Fig. 5(c) and (d). This particular experimental design was chosen as it is typical for global atmospheric models to attempt to parameterize sub-grid scale behavior (e.g. for phenomena occurring on a finer temporal/spatial scale).




(a) Two-dimensional view of orbit segment of system attractor.



(b) Three-dimensional view of orbit segment of system attractor.



(c) Two-dimensional view of model attractor.



(d) Three-dimensional view of model attractor.

Fig. 5. System (a–b) and model (c–d) attractor. Panels (a) and (c) show two-dimensional views looking at x_1 vs x_3 . Panels (b) and (d) show three-dimensional views using x_1 , x_2 , and x_4 ($I = 4, J = 16$).

3.1. Ensemble creation

The experimental design is inspired by and follows that of Wilks [26]. First, long integrations of the system on randomized initial values were performed to establish the shape of the attractor. A set of 500 different I -dimensional points were then chosen at 250-day intervals. This spacing was chosen to ensure that the initial states sampled different regions of the system attractor, and neighboring states were uncorrelated. A ‘true’ trajectory \mathbf{z}^a from each of these states was determined using a 50-day integration of the system. The goal is to use the model to shadow each of these individual true trajectories with an ensemble of 20 members. Care was taken to establish an ensemble that was a dynamically realistic sample of state space near the initial state as follows.

At each of the 500 initial states, an I -dimensional hypersphere was constructed encompassing 100 neighboring states from the system attractor, identified during a very long system integration. A neighboring state is defined to be one within 5% of the span of x_i in the i th dimension of the attractor. For each hypersphere, the covariance of the 100 neighboring states is calculated, yielding a $I \times I$ matrix C . The matrix values are then scaled to ensure that the average ensemble standard deviation is 5% of the climatological span of the system attractor. Thus, for each of the 500 hyperspheres, 20 initial ensemble members are drawn from the distribution:

$$C^{\text{init}} = \frac{0.05^2 \tau^2}{\lambda} C, \quad (2)$$

where λ is the average eigenvalue of C and τ is the standard deviation of x_i observed in the attractor.

First, a control state is picked within each hypersphere by adding appropriately distributed random noise to the I slow variables of the truth as follows:

$$\mathbf{z}_0^f(i, 1) = \mathbf{z}_0^a(i, 1) + \sqrt{C^{\text{init}}}\mathbf{y}(i, 1), \quad (3)$$

for $i = 1, 2, \dots, I$. The vector \mathbf{y} is I -dimensional, consisting of random entries from a Gaussian distribution. The Cholesky decomposition is used to calculate the square root of C^{init} [32]. Note that if C^{init} was not symmetric positive definite (due to the finite sample in time), then the lower triangular matrix was matched to the upper triangular matrix. The remaining 19 ensemble members are then chosen using the same method, but using the control state $\mathbf{z}_0^f(i, 1)$ for $i = 1, 2, \dots, I$ as the central reference point. Thus, we have

$$\mathbf{z}_0^f(i, m) = \mathbf{z}_0^f(i, 1) + \sqrt{C^{\text{init}}}\mathbf{y}(i, 1), \quad (4)$$

for $i = 1, 2, \dots, I$ and $m = 2, 3, \dots, 20$, where \mathbf{z}_0^f is a matrix with ensemble members as columns. Finally, the initial fast variables for the forecast are set equal to the initial fast variables of the truth, $\mathbf{z}_0^f(j, m) = \mathbf{z}_0^a(j, 1)$ for $j = I + 1, I + 2, \dots, I(j + 1)$ and $m = 1, 2, \dots, 20$. These fast variables are subsequently ignored when measuring the shadowing time.

3.2. Characterizing the ensemble

Once the ensemble of initial states has been created for each hypersphere, the trajectory of each ensemble member is forecast




Fig. 6. A two-dimensional representation of a 20-member ensemble ($I = 2$). The black point is the control forecast, and the white points are the other ensemble members. Let A be the 2×19 matrix of ensemble members with the control forecast defined as the origin. The semi-axes are calculated using SVD, where s_1 and s_2 are the two singular values of the matrix A . Note that the SVD spectrum can be quite steep for forecasts longer than a few time steps.

using the model. Let A be the $I \times (m - 1)$ matrix of ensemble members with the control state at the origin, where m is the number of ensemble members (including the control forecast) and we are only considering slow variables. If the matrix A acts on every vector in the $(m - 1)$ -dimensional unit sphere, the result is a filled-in ellipsoid, provided $m > I$. The points within the ellipsoid can be thought of as representing all potential ensemble members based upon the known m members (all of which lie within this ellipsoid). Although the system is chaotic, over a short time span one can approximate the nonlinear evolution of the initial sphere of trajectories with this ellipsoid [33].

Every two time steps, the ellipsoid is analyzed to determine expanding and contracting directions (i.e. those directions with positive/negative local finite-time Lyapunov exponents). This frequency was chosen rather than every time step to ensure the identified directions continue to expand/contract. Let $A = USV^T$ be a singular value decomposition (SVD) of A , and let $\mathbf{s}(t) = s_1, s_2, \dots, s_I$ be the singular values (the diagonal entries of S) and $U(t) = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_I]$ the matrix of left-singular vectors at a given time t . Thus the lengths and directions of the semi-axes of the ellipsoid approximately encompassing all potential ensemble members are given by $s_i \mathbf{u}_i$ [32]. At each time interval, the vector $\mathbf{s}(t)$ is compared to $\mathbf{s}(t - 1)$ (the singular values from the previous time interval) to determine which singular values are decreasing. We then computed the dot product between all possible combinations of former $u_i(t - 1)$ and current $u_i(t)$ left singular vectors, with the largest dot product chosen to identify direction pairs. Note that on occasion, a perfect matching was not achieved, but the best guess was used. A two-dimensional analogue is shown in Fig. 6. Note all ensemble members lie within the ellipse, and the control state is located at the center by definition of A . SVD allows us to identify the ellipse semi-axes $s_1 \mathbf{u}_1$ and $s_2 \mathbf{u}_2$.

An example one-dimensional 40 day ensemble forecast is shown in Fig. 7 ($I = 6, J = 16$). Note that for the first 20 days, almost all ensemble members remain close to the truth. However, by day 40 the spread of ensemble members covers the entire state space, and the ensemble mean no longer accurately represents the truth. In addition, Fig. 7 illustrates that the model achieves greater extreme values than the system. With less energy being dissipated through the fast variables, more energy remains oscillating between the slow variables in the model, as compared to the system. Note that the reduced regularity of the model relative to the system can also be observed.




Fig. 7. A 20-member ensemble forecast for x_3 initialized using Eq. (4). The black line represents the truth, an integration of the system. The gray lines are individual ensemble members, integrations of the model ($I = 6, J = 16$).




Fig. 8. Schematic of forecasting in two dimensions. The black line represents the trajectory of the truth (\mathbf{z}^a). \mathbf{Z}_t^f is the ellipsoid encompassing the ensemble members at time t , and the dashed line is the trajectory of the ensemble mean. For nonlinear systems, this ellipsoid is only a good approximation of the collections of forecasts for small t . Note that the time steps 0, 1, 2, 3 are an incomplete sampling of all time steps. The goal of a forecast is to achieve an ensemble which has a nonempty intersection with \mathbf{z}_F^a . Here forecasting fails after time 2.

4. Forecasting

Forecasting often fails when the state space dynamics of the physical system expand along a previously contracting dimension. A two-dimensional illustration of this concept is shown in Fig. 8. Because of the uncertainty associated with the initial location of the truth, an I -dimensional sphere of radius σ is used to represent the truth. It is from within this sphere that the initial ensemble members are chosen. The ellipse representing the ensemble members diverges from the trajectory of the truth when the system quickly turns in a new direction. In the illustration, forecasting has failed at time F as there is no overlap between the I -dimensional sphere around the truth \mathbf{z}_F^a and the ensemble ellipsoid \mathbf{Z}_F^f . For all experiments, the number of expanding directions observed throughout the forecast was non-constant. This unstable dimension variability is a cause of shadowing failures well-documented in the literature [22,23,34,35].

4.1. Forecasting with inflation

As described in the introduction, inflation is performed in an attempt to improve the ensemble forecast by expanding the ellipsoid along contracting dimensions of the state space. For each direction \mathbf{u}_i in which the ensemble is contracting, ensemble members are inflated along the i th semi-minor axis of the ellipsoid $s_i \mathbf{u}_i$. Although this adds artificial uncertainty to the forecast, Danforth and Yorke argue that the amount is minimal [22]. If the




Fig. 9. Schematic of forecasting with inflation in two dimensions. The black line represents the trajectory of the truth. Z_i^f is the ellipsoid encompassing the ensemble members at time t , and the dashed line is the trajectory of the ensemble mean. At each time step, the ellipsoid is inflated by φ along contracting directions, where in this case there is only one such direction. Note that forecasting now achieves its goal in contrast to Fig. 8.

ellipsoid continues to contract in these dimensions, the synthetic uncertainty continuously decreases. However, if the system begins to expand along a previously contracting dimension, the inflated ensemble will capture some of the change. This concept is illustrated for two dimensions in Fig. 9.

As described above, the directions of the semi-axes are calculated using SVD throughout the forecast at periodic time intervals and $\mathbf{s}(t)$ is compared to $\mathbf{s}(t - 1)$. For each i in which $s_i < s_i(t - 1)$, all ensemble members are inflated in direction \mathbf{u}_i . It is important to note here that if inflation were performed in expanding dimensions as well, forecasts would degrade more rapidly. By constraining inflation to the contracting dimensions, we are minimizing the damage caused in hyperbolic regions of state space.

Let φ be the inflation amount, expressed as a percentage of the semi-axis length, and let A be the matrix of ensemble members with the control state at the origin, prior to inflation. Define

$$M = \mathbb{I} + \varphi(\mathbf{u}_i \mathbf{u}_i^T) \quad (5)$$

where \mathbb{I} is the I -dimensional identity matrix, \mathbf{u}_i^T is the transpose of \mathbf{u}_i , and let $A' = MA$. Adding the control state to each vector in A' then yields ensemble members inflated by φ in direction \mathbf{u}_i .

The benefit of inflation for a particular forecast is illustrated in Fig. 10 for a two-dimensional projection of an ensemble forecast ($I = 6$). As the trajectory of \mathbf{z}^a begins to turn clockwise, the forecasts (gray lines) diverge from the black line. As represented by the gray lines, forecasting fails at location 1 without inflation

(Fig. 10(a)), when there is no longer overlap between the σ -sphere surrounding \mathbf{z}^a and the ensemble ellipsoid. However, in Fig. 10(b), the gray lines are inflated at location 1 in previously contracting directions (as inferred from SVD of the model ensemble). Note that the gaps in the figure represent both an inflation and a time step. As the figure illustrates, these shifts help capture the change in direction for \mathbf{z}^a . It is important to note that no measurements of \mathbf{z}^a are used to determine the directions in which to inflate. A second inflation for the gray lines occurs four time steps later at location 2 in panel b. At first there is more uncertainty in the ellipsoid encompassing the ensemble members. However, as \mathbf{z}^a continues in the same direction, the uncertainty is dampened by the contracting dynamics. Note that although the ellipsoid is analyzed every two time steps, the model did not inflate between locations 1 and 2. Also, inflation acts symmetrically on an ensemble, leaving the ensemble mean unchanged during a single step. However, once the ensemble is integrated forward in time, the inflated and non-inflated ensemble mean forecasts diverge.

4.2. Results

For each of the 500 initial hyperspheres, a 50-day model forecast for the 20 ensemble members was created with and without inflation. At each point in time, let \mathbf{z}^* be the mean of the 20 ensemble members' slow variables. For each hypersphere, \mathbf{z}^* was compared to the slow variables of the truth, where we let \mathbf{z}^a represent only the slow variables for simplicity. The root mean square error (RMSE) and anomaly correlation (AC) were calculated at each time step as follows:

$$\text{RMSE} = \|\mathbf{z}^* - \mathbf{z}^a\|_2 \quad (6)$$

$$\text{AC} = \frac{(\mathbf{z}^* - \bar{\mathbf{z}}) \cdot (\mathbf{z}^a - \bar{\mathbf{z}})}{\|\mathbf{z}^* - \bar{\mathbf{z}}\|_2 \|\mathbf{z}^a - \bar{\mathbf{z}}\|_2} \quad (7)$$

where $\bar{\mathbf{z}}$ represents the vector of system climatological averages for each slow variable. RMSE and AC calculations were then averaged over all hyperspheres for the duration of each forecast. For the system, $\text{RMSE} \approx 9$ represents saturation, at which point there is no similarity between the forecast and the truth.

For the present discussion, we consider the period for which a forecast is useful, namely the time for which AC remains above 0.6 [7]. For each state space size, the average useful time was calculated using the non-inflated forecasts. This yielded durations of 6.7, 8.7, and 17.0 days, corresponding to $I = 4, 5$, and 6,

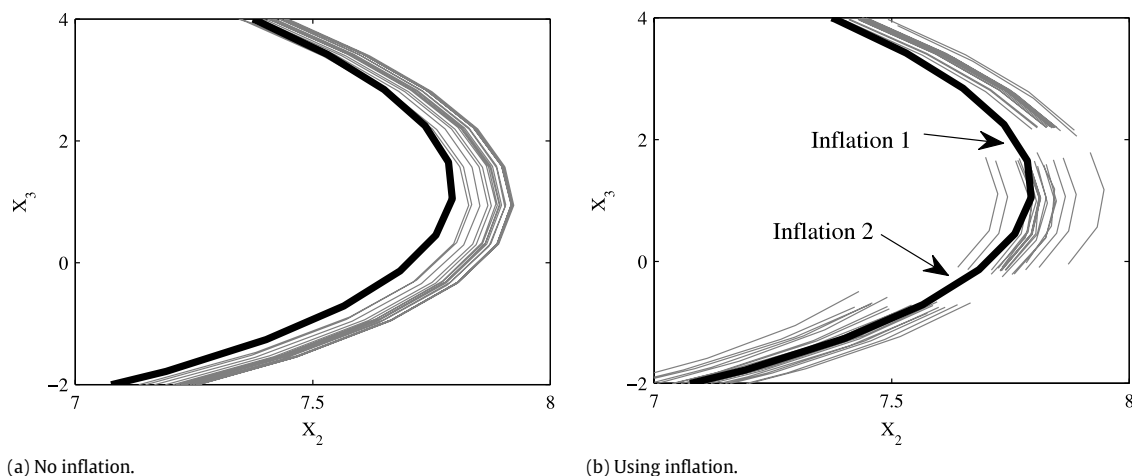


Fig. 10. Successful inflation in an ensemble forecast. The black line is the true trajectory \mathbf{z}^a , and the gray lines are ensemble member forecasts without inflation (panel a) and with inflation having the potential to occur at a minimum of every two time steps (panel b). All trajectories are moving in a clockwise direction. At location 1, forecasting fails without inflation. Location 2 represents a second inflation of the ensemble. Inflation successfully modifies the ensemble to include the truth as a plausible member. Note that no observations of \mathbf{z}^a are used during the inflation at locations 1 and 2 ($I = 6, J = 16$).




Fig. 11. Inflation forecasting results for 500 hyperspheres. Bars above zero show the number of hyperspheres for which inflation improved the forecast, while bars below zero show the number of hyperspheres for which inflation degraded the forecasts. Black bars indicate improvements and degradations of the average shadowing time by more than 5% (“inflation succeeded” and “inflation failed”, respectively). Gray bars represent the categories “inflation helped” and “inflation hurt”. The number of times inflation improved the forecast generally exceeded the number of degradations, both of which tend to increase with dimensionality and φ .

respectively. If the inflated and non-inflated forecasts crossed the 0.6 threshold during the same time step, the useful times were considered indistinguishable. Inflation was deemed to have “succeeded” if the duration for an acceptable forecast improved by more than 5% of the average useful time for $\varphi = 0$ forecasts. Similarly, the technique was considered to have “failed” if the duration worsened by the same threshold. If there was any improvement (greater than one time step) with inflation, forecasts were denoted “inflation helped”. Any measurable degradation with inflation was labeled “inflation hurt”. Results for $I = 4, 5$, and 6, with $\varphi = 0.5\%, 1\%, 2\%$, and 5%, are recorded in Fig. 11 with totals out of 500. Bars above zero represent improvements, while bars below zero represent degradations. Black bars indicate improvements and degradations of more than 5% of the average useful time (i.e. the number of successes and failures). Other metrics such as RMSE and the trajectory of the best ensemble member were considered, but they showed the same general trends. For all experiments described by Fig. 11, $J = 16$ as illustrated in Fig. 3.


Although the sample size is small, Fig. 11 suggests that inflation has a greater effect for larger I , with the best results exhibited

for $I = 6$ and $\varphi = 1\%$. In fact, the number of hyperspheres in which inflation helped was an order of magnitude better for $I = 5$ than for $I = 4$. The vast majority of inflated $I = 4$ forecasts were identical to non-inflated ($\varphi = 0$) forecasts. These numbers roughly doubled for $I = 6$ relative to $I = 5$, likely due to the increased regularity in \mathbf{z}^a for $I = 6$ observed in Fig. 3. Unfortunately, the same trend is seen for forecasts in which inflation failed. By capturing one of the sudden changes in the trajectory of the truth using inflation, the model can better forecast the truth for a longer duration. Alternatively, this increase in the occurrence of inflation helping might be evidence that error introduced by inflation in contracting directions has a minimal effect, as argued by Danforth and Yorke [22]. This effect would likely be strongest in systems with a steep singular value spectrum, where the expanding dimensions play a more dominant role.

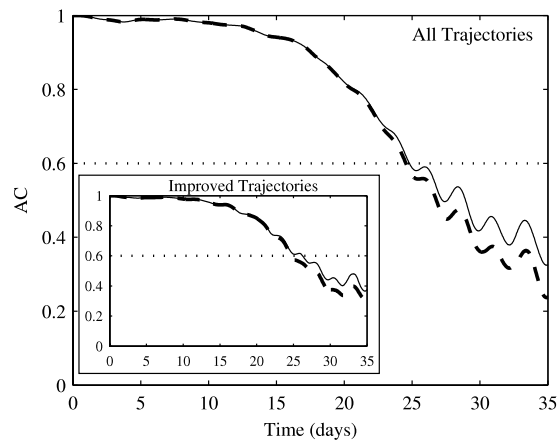
Naturally, the success of inflation is also strongly dependent on φ , with increasing φ generally corresponding to a larger number of both successes and failures. By increasing the inflation amount, the ellipsoid encompassing the ensemble members is more likely to overlap with the trajectory of the truth, and thus capture unstable dimension variability events. On the other hand, inflation can make a forecast worse. By inflating in directions incommensurate with the truth, the method introduces additional error, which can be quite significant when amplified by nonlinearity. The number of hyperspheres in which inflation failed also naturally increases with φ . However, as mentioned above, provided the contracting directions continue to contract, the error introduced is hopefully minimal.

For a given set of experimental parameters, the averaged useful time over all 500 hyperspheres is used for analysis. However, the $\varphi = 0$ forecast provides an individual useful time for each hypersphere. Whether or not inflation helps or hurts (i.e. reduces or increases the individual useful time) is not dependent upon the magnitude of the individual useful time. For example, inflation is equally likely to help (or hurt) an individual forecast whose non-inflated useful time was 3 days or 20 days.

As shown in Fig. 11, for many of the hyperspheres inflation successfully improved the duration for which a forecast is useful. The lack of harm caused by this technique in a global sense can be seen by averaging RMSE and AC over the hyperspheres in which inflation was successful as shown in Fig. 12 for $I = 6, J = 16$. Both plots show small improvement relative to forecasts made with no inflation. In Section 6, we demonstrate the ability to isolate conditions likely to result in successful inflation as in Fig. 10.



(a) RMSE.



(b) AC.

Fig. 12. Averaged RMSE (a) and AC (b) for the $I = 6, J = 16$ system. The dashed line represents no inflation, while the solid line is with inflation of 1%. Figures show the averaged data for all 500 hyperspheres, while insets only include the hyperspheres for which inflation succeeded. Given the small (positive) change in average forecast statistics, the method is unlikely to harm forecasts, and may be adjusted to isolate particular events for which inflation can be quite successful (see Section 6).

Conducting experiments for larger dimensional systems was difficult due to computation time limitations for estimating the local dynamics of the system attractor (see Section 3). To reduce the computational burden, initial ensemble members were chosen using a spherically-symmetric Monte Carlo approach (i.e. nearby initial conditions but with no attractor-imposed structure) for these larger systems ($I = 7, 8, \dots, 40$). The $I = 6, I = 16$, and $I = 21$ systems had substantially greater average useful times than the other systems (by a factor of five). These specific systems all exhibit increased regularity in their time series (see Fig. 3), presumably a result of some resonance between the slow and fast variables. For these systems the model is most sensitive to inflation, with small φ greatly improving forecasts and larger φ greatly degrading forecasts. The less-regular systems had smaller incidence of success and failure, with results similar to those observed for the $I = 4$ and $I = 5$ systems. However, the use of a spherically-symmetric Monte Carlo approach to initialize an ensemble degraded the performance of inflation when compared to the more sophisticated ensemble initialization technique.

5. Shadowing and stalking

5.1. Methods

Shadowing with an ensemble presents an alternative method for assessing the quality of a model's predictions. Contrasting the forecasting experiments described previously, in this context the trajectory of the truth is known a priori to within a given uncertainty (σ), and its location at a given time is represented by an I -dimensional sphere of radius σ . At regular intervals, the intersection between the ellipsoid encompassing the ensemble members and the σ -sphere is approximated. The ensemble ellipsoid is then redefined to represent the intersection, while the other ensemble members are redefined to fall within the overlap. Provided this intersection is nonempty, the trajectories of some ensemble members yield accurate representations of reality. A schematic of this technique is shown in Fig. 13. Note that although the actual truth lies outside the ensemble ellipsoid, some ensemble members are within σ . The trajectories of these ensemble members can be forecast forward in time, and the process is repeated. If some ensemble members remain in the intersection for the entire forecast, then their trajectories have successfully σ -shadowed the truth. When inflation is applied to the ensemble ellipsoid (as described previously), this aggressive form of shadowing is called *stalking*.

5.2. Results

The same initial conditions utilized in the forecasting experiments were used to create 50-day stalking ensembles. However, every four time steps the ensemble ellipsoid was redefined as described in Fig. 13. This frequency was chosen to be long enough to ensure that appropriate contracting dimensions could be identified. During the definition of the overlap between the ensemble ellipsoid and the truth sphere, semi-axes can be altered drastically, requiring several time steps to reflect the local dynamics. The σ -shadowing distance (uncertainty around the truth) was chosen to be 10% of the climatological range in each dimension, namely twice the spread in the initial ensemble. A new collection of 20 ensemble members were then chosen lying within the redefined ellipsoid. As before, RMSE and AC were calculated and compared for both shadowing and stalking experiments. Counts for the number of hyperspheres for which inflation succeeded, failed, helped, and hurt as described in Section 4 were tallied. Results for $I = 4, 5$, and 6, with $\varphi = 0.5\%, 1\%, 2\%$, and 5% , are given in Fig. 14.




Fig. 13. Schematic of the redefining of a shadowing ensemble in two dimensions. The star represents the known location of the truth. The solid circles are supplemented by a collection of points meant to approximate the overlap of the ellipse and the sphere around the truth. The open circles are previous ensemble members that will now be redefined within the overlap, which is approximated by a new ellipsoid. Note that the x_1, x_2 axes are exaggerated for visual ease.




Fig. 14. Stalking results for 500 hyperspheres. Bar categories are identical to those in Fig. 11. Inflation improves and worsens the ability to shadow with similar probability. Counts tend to slightly increase with the dimension of the system.

The shadowing experiments exhibited a greater benefit from inflation when the dimensionality I is large, although the difference is less than before. However, the number of hyperspheres for which inflation succeeded does not seem to trend monotonically with φ over the range $0.5\%–5\%$. Since the number of hyperspheres in which inflation succeeded roughly equals the number for which inflation helped, we can conclude that when stalking is constructive, it improves the shadowing time by more than one day. On the other hand, the number of hyperspheres in which inflation succeeded is nearly equivalent to the number for which inflation failed for all experiments. Thus, inflation can also have a negative effect upon shadowing.

As with forecasting, stalking experiments were extended to higher dimensional systems using a spherically-symmetric Monte Carlo approach for ensemble initialization. Similar to the lower dimensional systems in Fig. 14, counts for the number of hyperspheres for which inflation succeeded and failed were relatively close. For small values of inflation, there were often errors in the procedure by which corresponding ensemble semi-axes directions are calculated between time steps, resulting in an increase in the number of failures. With larger inflation, the algorithm can better determine corresponding directions during redefinition of the ensemble. The highly regular $I = 6, I = 16$, and

$I = 21$ systems were less distinguished for stalking experiments. Since the truth is known throughout, ensemble members are forced by the redefinition to remain nearby regardless of the regularity of the attractor.

6. Targeted inflation

Despite taking advantage of local stability estimates, inflation as defined previously was naive in the sense that it occurred in all contracting directions, including those which push forecasts away from the truth. Thus, this naive inflation had the potential to both improve and degrade a forecast. The final objective of this research was to isolate the instances in which inflation is beneficial to the forecast, and identify the systematic tendencies of these situations. With this knowledge, at each time step the decision of whether inflation should occur, and in which contracting directions, could be based upon the characteristics of the local state space of the system. For these *targeted inflation* experiments, the $I = 5$ system was chosen for further study; it does not have the same regularity as the $I = 6$ system (Fig. 3), but is more regular than some of the higher dimensional systems.

For a given time step in the forecast, the shape of the attractor near the current state \mathbf{z}_t^f was approximated as follows. First, the state 50 time steps prior in the forecast \mathbf{z}_{t-50}^f was considered. By adding uniform random noise in every direction on the order of σ , 1000 analogues located near (within σ) \mathbf{z}_{t-50}^f were chosen. These 1000 analogues were forecast forward 50 time steps, and the covariance matrix associated with \mathbf{z}_t^f was calculated. The eigenvectors of this forecast state covariance matrix yield an approximation of the shape of the attractor near \mathbf{z}_t^f . Finally, the dot product of these eigenvectors and the proposed inflation direction (determined by SVD to be contracting) was calculated. The larger this projection, the more likely the proposed inflation direction is consistent with the local shape of the attractor. If the resulting dot product was above a set threshold μ and the direction was determined to be contracting, then inflation proceeded as before. Otherwise, no inflation occurred. For all values of ϕ , μ was set at 0, 0.8, and 0.9, reducing the number of inflations by roughly 0%, 60%, and 80%, respectively.

The goal here is to avoid inflating in contracting directions likely to harm the forecast. By initializing the analogues randomly, we are able to sample a new collection of state space directions, distinct from those spanned by the inflated ensemble.

6.1. Forecasting

Forecasting results are given in Fig. 15. Adjusting the number of inflations throughout the forecast by changing μ clearly improves the method. Note first that results from a cutoff of $\mu = 0$ agree with Fig. 11. The small differences (most notable for $\phi = 1\%$) can be attributed to starting from different initial points in the attractor. For $\phi = 2\%$ and $\phi = 5\%$, increasing μ has the effect of decreasing the counts; with fewer inflations, ensembles more closely resemble the non-inflated forecast ($\phi = 0$). It is important to note that although the number of hyperspheres for which inflation succeeds remains relatively constant, the number for which inflation fails is substantially reduced as μ is increased. This trend indicates that the technique is successfully determining when inflation will indeed be useful, and is inflating less frequently when it is likely to worsen a forecast. For example, consider $\phi = 5\%$ in Fig. 15. With naive inflation ($\mu = 0$), the technique succeeded for 39 hyperspheres and failed 42 times. Setting $\mu = 0.9$ reduced the number of successes by 23%–30%, however, it eliminated 71% of the failures, leaving only 12. When μ is increased beyond 0.9




Fig. 15. Targeted inflation for forecasting experiments on 500 hyperspheres. Bar categories are the same as those defined in Fig. 11. Note that $\mu = 0.9$ substantially reduces the instances of failure, while not significantly reducing the instances of success ($I = 5$).




Fig. 16. Targeted inflation for stalking experiments on 500 hyperspheres. Bar categories are the same as those defined in Fig. 11. Results varied, with more work needed to optimize the success of inflation for stalking experiments ($I = 5$).

though, inflation occurs too rarely and the number of hyperspheres for which inflation succeeds is greatly reduced (not shown).

For $\phi = 0.5\%$ and $\phi = 1\%$, the trend is a little more complicated. Still, $\mu = 0.9$ resulted in a remarkably better ratio of number of hyperspheres succeeding versus failing. With $\phi = 0.5\%$, $\mu = 0.8$ surprisingly increased the number of hyperspheres succeeding and helping relative to $\mu = 0$. For $\phi = 1\%$, $\mu = 0.8$ had notably few hyperspheres for which inflation succeeded. This indicates that although experiments were run for 500 different hyperspheres, there is still variability between trials based upon initial conditions and chaotic system dynamics.

6.2. Stalking

As shown in Fig. 16, the results for stalking experiments are more variable. For some algorithm parameters, such as $\mu = 0.8$ for $\phi = 0.5\%$ and $\phi = 1\%$, targeted inflation appears to successfully increase the number of hyperspheres for which inflation succeeded, while decreasing the number for which inflation failed. However, the same μ for greater inflation had smaller differences, even increasing the number of trajectories for which inflation failed. One possible explanation for the mixed results in stalking experiments is the experimental design. Redefining the ensemble every fourth time step reduces the ability to identify expanding/contracting dimensions. This can lead to

difficulties in matching ellipsoid semi-axis directions from one time step to the next. Adjusting μ did alter the results, however more work is needed to optimize the success of inflation for stalking experiments.

7. Discussion

Using the Lorenz '96 coupled system as an analogue for atmospheric dynamics, we examined the potential for using inflation to improve model forecasts of the Earth's atmosphere. In the first experiments, forecasts were made with and without inflation, and the resulting trajectories were compared to the system truth. In the second part of this work, the existence of a shadowing trajectory with and without inflation was assessed. The trajectory of the truth was known a priori, and only the closest ensemble members were considered every few time steps. In the former, the technique is referred to as inflation, and in the latter, stalking. Under idealized conditions, inflation was shown to be beneficial for both techniques. The final part of the work attempted to isolate the ensemble directions most likely to benefit from increased uncertainty using targeted inflation.

For the Lorenz '96 model, by adjusting the number of slow and fast variables, and the degree of coupling between them, one could vary between regular and highly chaotic systems; the greatest regularity was observed for $l = 6$ and $l = 16$. However, by increasing the number of fast variables (J), greater regularity can be achieved for a fixed number of slow variables. In contrast to hyperbolic systems, the present system exhibits unstable dimension variability and fluctuating Lyapunov exponents as is common in higher dimensional systems [35].

For forecasting experiments, inflation showed the potential to increase the length of time for which an ensemble forecast represents the verifying truth as a plausible member, as illustrated in Fig. 10. Inflation in directions away from the truth, however, can degrade a forecast. As a result, greater inflation amounts always correspond to an increase in the number of hyperspheres for which inflation failed. For regular systems such as $l = 6$ and $l = 16$, smaller inflation amounts can often successfully improve a forecast, while the error introduced from larger inflation amounts is magnified. However, for less regular systems, larger inflation was necessary to improve forecasts. For these systems, the uncertainty introduced was less likely to have a harmful effect.

Stalking experiments were used to establish the existence of a model trajectory that remains close to the truth for a given time, and assess how inflation affects this trajectory. For these experiments, results were relatively constant for systems with greater than six dimensions. Inflation was shown to be successful at about the same rate at which it hindered shadowing. Thus, the utility of stalking is dependent upon the ability to isolate the hyperspheres for which inflation helped.

Finally, by quantifying the local dynamics at each time step and restricting inflation to directions consistent with the local attractor, the instances for which inflation helped were isolated with moderate success for forecasting experiments. The number of hyperspheres for which inflation failed could be decreased without significantly decreasing the number for which inflation succeeded. Stalking experiments with targeted inflation yielded mixed results, with additional modifications needed to attain the same degree of success. In practice however, the truth is not known a priori. Thus, the improvements observed with targeted inflation for forecasting experiments are encouraging for application in NWP.

When shadowing physical systems such as the Earth's atmosphere, one is presented with the challenges of uncertainty in the initial condition, sensitive dependence, and model error. Model error is estimated to dominate the forecast error in weather systems for the first three days [17]. Nevertheless, modelers must attempt

to provide the best predictions possible given the circumstances. This research demonstrates that inflation has the potential to aid in our attempts to model chaotic physical systems, provided we are able to isolate the local shape of the attractor and tune the inflation parameters successfully. Targeted inflation is most likely to be beneficial for highly nonlinear events (e.g. hurricanes), where sudden changes in the state space trajectory are prevalent. Further exploration of the design parameters of targeted inflation are likely to result in increased effectiveness.

One potential application of this technique to NWP is in maintaining the spread of a forecast ensemble. For unbiased ensembles, an effective spread is equal to the RMSE of the ensemble mean with respect to the analysis [36], but NWP ensembles typically fail to retain this ideal beyond the first few days. Of the various means by which one could increase the ensemble spread artificially, targeted inflation preserves the ensemble orientation relative to the local model attractor and has the potential to improve the ensemble mean.

The major disadvantage of targeted inflation, relative to the naive version, is the frequently required computation of the shape of the attractor. For NWP models, where the number of degrees of freedom is $O(10^{10})$, this computation would be prohibitive. However, the operational forecasting centers compute new analyses and ensemble forecasts roughly every six hours. Thus, the required information for an estimate of the attractor shape would be available cost free in terms of computational time.

Acknowledgments

All experiments were performed using the Vermont Advanced Computing Center (VACC) 1400 processor cluster, an IBM e1350 High Performance Computing system. This research was supported by a National Aeronautics and Space Administration (NASA) EPSCoR grant and National Science Foundation (NSF) Grant # DMS-0940271. The original matlab code can be found at: <http://www.uvm.edu/~cdanfort/nolink/atmos/stalking-code.tar>. The authors would like to thank Kameron Harris, Nicholas Allgaier, and two anonymous reviewers for insightful comments and suggestions in preparing this manuscript.

References

- [1] C.E. Leith, Objective methods for weather prediction, *Annu. Rev. Fluid Mech.* 10 (1978) 107–128.
- [2] J.S. Whitaker, T.M. Hamill, X. Wei, Y. Song, Z. Toth, Ensemble data assimilation with the NCEP global forecast system, *Mon. Weather Rev.* 136 (2008) 463–482.
- [3] T.M. Hamill, J.S. Whitaker, Accounting for the error due to unresolved scales in ensemble data assimilation: a comparison of different approaches, *Mon. Weather Rev.* 133 (2005) 3132–3147.
- [4] T.M. Hamill, Ensemble-based atmospheric data assimilation, in: *Predictability of Weather and Climate*, 2006, pp. 124–156.
- [5] B.R. Hunt, E.J. Kostelich, I. Szunyogh, Efficient data assimilation for spatiotemporal chaos: a local ensemble transform kalman filter, *Physica D* 230 (2007) 112–126.
- [6] E. Ott, B.R. Hunt, I. Szunyogh, A.V. Zimin, E.J. Kostelich, M. Corazza, E. Kalnay, D.J. Patil, J.A. Yorke, A local ensemble kalman filter for atmospheric data assimilation, *Tellus A* 56 (2004) 415–428.
- [7] E. Kalnay, *Atmospheric Modeling, Data Assimilation and Predictability*, Cambridge Univ. Press, 2003.
- [8] E.N. Lorenz, A study of predictability of a 28-variable atmospheric model, *Tellus* 17 (1965) 321–333.
- [9] Z. Toth, E. Kalnay, Ensemble forecasting at NMC: the generation of perturbations, *Bull. Am. Meteorol. Soc.* 74 (1993) 2317–2330.
- [10] D.V. Anosov, Geodesic flows on closed Riemannian manifolds with negative curvature, *Trudy Mat. Inst. Steklov* 90 (1967) 3–210.
- [11] R. Bowen, Ω -limit sets for axiom A diffeomorphisms, *J. Differential Equations* 18 (1975) 333–339.
- [12] C. Robinson, Stability theorems and hyperbolicity in dynamical systems, *Rocky Mountain J. Math.* 7 (1977) 425–437.
- [13] K.R. Meyer, G.R. Sell, An analytic proof of the shadowing lemma, *Funkcial. Ekvac.* 30 (1987) 127–133.
- [14] T. Sauer, J.A. Yorke, Rigorous verification of trajectories for the computer simulation of dynamical systems, *Nonlinearity* 4 (1991) 961–979.
- [15] K. Judd, Shadowing pseudo-orbits and gradient descent noise reduction, *J. Nonlinear Sci.* 18 (2008) 57–74.

- [16] L.A. Smith, C. Ziehmann, K. Fraedrich, Uncertainty dynamics and predictability in chaotic systems, *Q. J. R. Meteorol. Soc.* 125 (1999) 2855–2886.
- [17] D. Orrell, L. Smith, J. Barkmeijer, T.N. Palmer, Model error in weather forecasting, *Nonlinear Processes Geophys.* 8 (2001) 357–371.
- [18] B.A. Coomes, H. Kocak, K. Palmer, A shadowing theorem for ordinary differential equations, *Z. Angew. Math. Phys.* 46 (1995) 85–106.
- [19] S.M. Hammel, J.A. Yorke, C. Grebogi, Do numerical orbits of chaotic dynamical processes represent true orbits? *J. Complexity* 3 (1987) 136–145.
- [20] C. Grebogi, S.M. Hammel, J.A. Yorke, T. Sauer, Shadowing of physical trajectories in chaotic dynamics: containment and refinement, *Phys. Rev. Lett.* 65 (1990) 1527–1530.
- [21] S.M. Hammel, J.A. Yorke, C. Grebogi, Numerical orbits of chaotic processes represent true orbits, *Bull. Amer. Math. Soc.* 19 (1988) 465–469.
- [22] C.M. Danforth, J.A. Yorke, Making forecasts for chaotic physical processes, *Phys. Rev. Lett.* 96 (2006) 144102.
- [23] S. Dawson, C. Grebogi, T. Sauer, J. Yorke, Obstructions to shadowing when a Lyapunov exponent fluctuates about zero, *Phys. Rev. Lett.* 73 (1994) 1927–1930.
- [24] H. Li, E. Kalnay, T. Miyoshi, Simultaneous estimation of covariance inflation and observation errors within an ensemble kalman filter, *Q. J. R. Meteorol. Soc.* 135 (2009) 523–533.
- [25] E.N. Lorenz, Predictability: a problem partly solved, in: *Proc. ECMWF Seminar on Predictability*, vol. 1, 1996, pp. 1–18.
- [26] D.S. Wilks, Effects of stochastic parametrizations in the Lorenz'96 system, *Q. J. R. Meteorol. Soc.* 131 (2005) 389–407.
- [27] D. Orrell, Role of the metric in forecast error growth: how chaotic is the weather? *Tellus A* 54 (2002) 350–362.
- [28] C.M. Danforth, E. Kalnay, Using singular value decomposition to parameterize state-dependent model errors, *J. Atmospheric Sci.* 65 (2008) 1467–1478.
- [29] E.N. Lorenz, K.A. Emanuel, Optimal sites for supplementary weather observations: simulation with a small model, *J. Atmospheric Sci.* 55 (1998) 399–414.
- [30] W.B. Hayes, A.V. Malykh, C.M. Danforth, The interplay of chaos between the terrestrial and giant planets, *Mon. Not. R. Astron. Soc.* 407 (2010) 1859–1865.
- [31] A.J. Simmons, R. Mureau, T. Petroliajgis, Error growth estimates of predictability from the ECMWF forecasting system, *Q. J. R. Meteorol. Soc.* 121 (1995) 1739–1771.
- [32] G. Golub, C.F. Van Loan, *Matrix Computations*, third ed., The John Hopkins University Press, 1996.
- [33] K. Alligood, T. Sauer, J.A. Yorke, *Chaos: An Introduction to Dynamical Systems*, Springer-Verlag, 1996.
- [34] G.-C. Yuan, J.A. Yorke, An open set of maps for which every point is absolutely nonshadowable, *Proc. Amer. Math. Soc.* 128 (2000) 909–918.
- [35] T. Sauer, Shadowing breakdown and large errors in dynamical simulations of physical systems, *Phys. Rev. E* 65 (2002) 36220.
- [36] D.S. Wilks, On the reliability of the rank histogram, *Mon. Weather Rev.* 139 (2011) 311–316.