AGGRESSIVE SHADOWING OF A LOW-DIMENSIONAL MODEL OF ATMOSPHERIC DYNAMICS

A Thesis Presented

by

Ross Lieb-Lappen

 to

The Faculty of the Graduate College

of

The University of Vermont

In Partial Fulfillment of the Requirements for the Degree of Master of Science Specializing in Mathematics

October, 2010

Accepted by the Faculty of the Graduate College, The University of Vermont, in partial fulfillment of the requirements for the degree of Master of Science, specializing in Mathematics.

Thesis Examination Committee:

_____ Advisor

Christopher Danforth, Ph.D.

Jun Yu, Ph.D.

_____ Chairperson

Joshua Bongard, Ph.D.

_____ Dean, Graduate College

Domenico Grasso, Ph.D.

Date: May 18, 2010

ABSTRACT

Modeling Earth's atmospheric conditions is difficult due to the size of the system, and predictions of its future state suffer from the consequences of chaos. As a result, current weather forecast models quickly diverge from observations as uncertainty in the initial state is amplified by nonlinearity. One measure of the strength of a forecast is its *shadowing time*, the period for which the forecast is a reasonable The present work uses the Lorenz '96 coupled system, a description of reality. simplified nonlinear model of atmospheric conditions, to extend a recently developed technique for lengthening the shadowing time of a dynamical system. An ensemble of initial states, systematically perturbed using knowledge of the local dynamics, is used to make a forecast. The experiment is then repeated using *inflation*, whereby the ensemble is regularly expanded along dimensions whose uncertainty is contracting. The first goal of this work is to compare the two forecasts to reality, chosen to be an imperfect version of the same model, and determine whether variance inflation succeeds. The second goal is to establish whether inflation can increase the maximum shadowing time for a single member of the ensemble. In the second experiment the trajectory of reality is known a priori, and only the closest ensemble members are considered at each time step. When inflation is introduced to this technique, it is called *stalking*. Variance inflation was shown to have the potential to be successful, with the extent dependent upon algorithm parameters (e.g. size of state space, inflation amount). Under idealized conditions, the technique was shown to improve forecasts over 50% of the time. Under these same conditions, stalking also exhibited the potential to be useful. When only the best ensemble members were considered at each time step, the known trajectory could be shadowed for an entire 50-day forecast 50-75% of the time. However, if inflation occurs in directions incommensurate with the true trajectory, inflation can actually reduce stalking times. Thus, utilized appropriately, inflation has the potential to improve predictions of the future state of atmospheric conditions, and possibly other physical systems.

ACKNOWLEDGMENTS

First of all, I would like to thank Chris Danforth for introducing me to mathematical research. When I felt lost or thought I had reached a dead-end, he had a new insight or different perspective that allowed me to continue. Although I started with hardly any programming experience, Chris was always incredibly patient with my ignorant questions. I may not be a LATEXdie-hard (yet), but I now appreciate its benefits. Our regular ping-pong games at least kept me sane.

I would like to thank Jim Lawson and the support staff for the Vermont Advanced Computing Center, who were always eager to answer my desperate pleas for programming assistance. Thank you to all of my professors in the Department of Mathematics and Statistics. Special thanks to Taras Lakoba for straightening my SVD questions while hiking on the Long Trail. To Jun Yu and Joshua Bongard, for serving on my committee and reading this work. Thank you to Nick Allgaier and Kameron Harris, for making our research group fun and enjoyable, and forcing me to transition to Linux.

Most importantly, thank you to all of my family and friends. Without your love and support, I would not be who I am today. Thank you mom, dad, Mia, and all of my grandparents for providing love and support regardless of the circumstances. Thank you Lee for staying with me in Vermont another two years as I try to find my academic path. You push me to reach my potential, while providing insight, inspiration, and encouragement along the way.

Acknowledgments ii
List of Tables iv
List of Figures v
Chapter
I. Introduction 1
II. Modeling 4
Dynamics
System vs. Model
III. Methods
Ensemble Creation
Making Predictions12
IV. Forecasting
Forecasting Results17
Monte Carlo Forecasting Results
V. Shadowing and Stalking
Stalking Results
Monte Carlo Stalking Results
VI. Conclusions
Appendix A: Understanding SVD
Appendix B: Matlab Code
Literature Cited

TABLE OF CONTENTS

LIST OF TABLES

4.1.	Variance inflation forecasting results out of 500 hyperspheres	21
5.2.	Shadowing and stalking results out of 500 hyperspheres	28

LIST OF FIGURES

2.1.	Model Schematic	4
2.2.	System perturbation time series	6
2.3.	Extended perturbation time series	6
2.4.	The dependence of coupling on system stability	7
2.5.	The dependence of the number of fast modes on system stability	8
2.6.	The dependence of the number of slow modes on system stability \ldots .	8
2.7.	A 10 day time series for x_1 and $y_1^{1,2,3,4}$ using $I = 8, J = 4 \dots$	9
3.8.	System Attractor	10
3.9.	Model Attractor	11
3.10.	A two-dimensional representation of a 20-member ensemble	13
3.11.	A 20-member ensemble forecast for x_3 initialized using Equation (5)	13
4.12.	Schematic of forecasting in two dimensions	14
4.13.	The number of expanding directions over a 50 day forecast	15
4.14.	Schematic of variance inflation in two dimensions	15
4.15.	Successful inflation in a forecast	16
4.16.	Representative averaged RMSE for different values of I	18
4.17.	Representative averaged AC for different values of $I \dots \dots \dots$	18
4.18.	Representative averaged RMSE for different inflation amounts φ	19
4.19.	Representative averaged AC for different inflation amounts φ	19
4.20.	Averaged RMSE and AC for hyperspheres in which inflation improved	
	forecast	22
4.21.	The average shadowing time for non-inflated forecasts using the Monte	
	Carlo approach for $I = 4$ to $I = 18$ dimensional systems	23
4.22.	Number of hyperspheres reaching the desired shadowing period of 10%	
	beyond the $\varphi=0$ average using the Monte Carlo approach $\ldots\ldots\ldots\ldots$	24

4.23. Number of hyperspheres where inflation worked and failed using the
Monte Carlo approach on $I = 4$ to $I = 18$ dimensional systems 2
5.24. Schematic of the redefining of an ensemble in two dimensions 2
5.25. Number of hyperspheres reaching the desired shadowing period of 10%
beyond average using the Monte Carlo approach 2
5.26. Number of hyperspheres where inflation worked and failed using the
Monte Carlo approach on $I = 4$ to $I = 18$ dimensional systems 3

I. INTRODUCTION

Society is often dependent upon the ability of scientists to accurately forecast the future state of chaotic physical systems, such as the weather. Meteorologists are asked to anticipate natural disasters with ample time for the public to adequately prepare. However, 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 [1]. In general, models fail to predict the behavior of chaotic physical systems due to uncertainty in the initial state, chaos, and model error [2]. The initial state is an estimate of the condition of the atmosphere at the start of a forecast. Although weather monitoring devices cover much of the planet, for many areas (e.g. oceans) the data is sparse. Therefore, meteorologists must estimate quantities in these regions using data assimilation, which introduces uncertainty to the forecast. The atmosphere is a chaotic system, meaning that these small uncertainties in the initial state are amplified by nonlinearity. Finally, since the model itself is not a perfect representation of reality, model error adds to the uncertainty leading forecasts to diverge quickly from observations.

Since perfect knowledge of the current state of the atmosphere 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, "the analysis". 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 data, the modeler's goal is to create a forecast that remains reasonable for the longest duration of time from this given state. In the example of weather, forecasters strive to be accurate for two weeks, the limit imposed by chaos. Today 5 day forecasts are as good as 3 day forecasts from 30 years ago [3]. To account for the initial state uncertainty, the accepted technique is to use ensemble forecasting, where a large collection of ensemble members are chosen randomly from

a distribution that reflects the system dynamics near the given initial state [4]. Each ensemble member is forecast forward in time, yielding a collection of final states. The probability distribution of this collection represents 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.

The time period for which a particular forecast is an accurate representation of reality is called the shadowing time. In 2006, Danforth and Yorke proposed a method called *stalking* to increase the shadowing time for a given forecast [5]. For a system with n-degrees of freedom, an n-dimensional disk is used to encompass the initial ensemble members. Throughout the length of the forecast, the system will be expanding in some dimensions, while it will be contracting in others, depending on the local finite time Lyapunov exponent in each dimension. The result can be approximated by an *n*-dimensional ellipsoid for a short time. The idea of stalking is to add some uncertainty in the contracting directions (i.e. those with a negative local finite time Lyapunov exponent) at periodic time steps 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 [6], some ensemble members will remain relatively close. This aggressive form of shadowing is known as stalking, and is not currently used in weather forecasting, but is used in 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.

Most of shadowing theory has been developed for hyperbolic systems based on the Shadowing Lemma of Anosov [7] and formalized by Bowen [8]. Given a pseudotrajectory of a model (i.e. one very close to an actual 1-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. [9–11]). 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, and there does not exist a trajectory that shadows the truth for an arbitrarily long time. More recent work has been focused on finding shadowing trajectories for non-hyperbolic systems (e.g. the driven pendulum, Hènon map) [12– 15]. In this paper, stalking is further developed using the non-hyperbolic '96 Lorenz system to evaluate the potential for improvement in forecasts [16].

II. MODELING

Ensemble forecasting is used to shadow a known trajectory \mathbf{z}^{a} , the "truth", of some meteorological quantity (such as temperature, pressure, etc.). 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. This system has been used in previous studies to illustrate weather related dynamics (e.g. [17–19]). The N-dimensional governing first-order differential equations are given by [16]:

$$\frac{dx_i}{dt} = x_{i-1}(x_{i+1} - x_{i-2}) - x_i + F - \frac{hc}{b} \sum_{j=J(i-1)+1}^{iJ} y_j \tag{1}$$

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

for i = 1, 2, ..., I and j = 1, 2, ..., JI, and N = (J + 1)I.



FIG. 2.1. Model Schematic. The I = 8 slow variables x_i can be thought of as locations along a latitude circle. Each variable has J = 4 corresponding y_i^j fast variables.

The values x_i represent slowly changing meteorological quantities whose dynamics are described by Equation (1). 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 Equation (2). For our first experiments, we set I = 4, 5, 6 and J = 16 for a state space of 68,81 or 102 variables. We then examined the effects of a larger state space by letting I vary from 4 to 18. A schematic is shown in Fig. 2.1. where 8 slow variables (x_i) are each coupled to 4 fast variables (y_i^j) . Note that the dynamics of each x_i are dictated by neighboring x variables and the corresponding set of coupled y_i^j variables.

The nonlinear terms in Equation (1) 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 h = 1, c = 10, and b = 10, which forces the fast variables to oscillate 10 times quicker than the slow variables [17, 19, 20]. Note that one time unit corresponds to 5 days, the dissipative decay time [21].

Dynamics

The system dynamics can be studied by considering integrations of Equations (1) and (2), hereafter referred to as the *system*. In Fig. 2.2, a time series for a longitudinal profile (I = 40, J = 16) is shown after x_{13} is initially perturbed by five units. Profiles are recorded at 12-h intervals over a 5-day forecast. Advection is apparent as the energy from the perturbation is observed halfway around the latitudinal circle by day 5. This initial energy pulse was allowed to propagate for 55 days. A time series for days 50-55 is shown in Fig. 2.3, and some of the same characteristics as seen in Fig. 2.2 can be observed. Lorenz and Emanuel [21] calculated a growth rate doubling time of approximately 2 days for this model, which agrees with trends in larger atmospheric models where errors double in roughly 2 days [22]. However, growth rates over limited



FIG. 2.2. System perturbation time series. The effect of a five unit perturbation is observed over a 5 day time series. After 5 days, the perturbation effects half of the locations, indicating advection to the east and more slowly to the west (I = 40, J = 16).



FIG. 2.3. Extended perturbation time series. The initial state of Fig. 2.2 is observed over a 5 day time series beginning at day 50. Particular perturbations can be traced throughout the time series (I = 40, J = 16).

time intervals as in Fig. 2.3 can differ greatly.

Adjusting the parameters of the system and examining the time series at a single site illustrates the dependence of the slow variables upon coupling to the fast variables. The relative significance of the fast modes can be observed by varying h as shown in Fig. 2.4 (I = 6, J = 16). For h = 1, a regular pattern emerges as an energy equilibrium is achieved between external forcing and dissipation. However, as the significance of the fast modes is reduced, the time series becomes more complex.



FIG. 2.4. The dependence of coupling on system stability. The time series for x_3 as the parameter h is varied. For all three panels, I = 6, J = 16.

Similar results can be achieved by changing the number of fast modes. As illustrated in Fig. 2.5, increasing J creates a more regular system. For I = 6, the system is fairly regular for both J = 16 and J = 40. However, with only a few fast modes active (e.g. J = 8), extreme oscillations become prevalent. For I = 8, more fast modes are required to achieve the consistent pattern. Note that for both J = 8and J = 16, the system is quite variable. Not until 40 fast modes are present does the system exhibit regularity. With more fast modes available, the system's energy is more evenly distributed and nonlinearities in the fast modes have less impact on the stability of the slow modes.

As shown in Fig. 2.6, adjusting the number of slow modes greatly alters the qualitative dynamics of the time series for each particular x_i . Note that for I = 4, 5 and 6, a regular pattern is observed with J = 16. As mentioned above, a consistent pattern could be attained for larger I by increasing J. However, the regular pattern varies drastically with I. 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 10 days. For I = 5, the system oscillates at an intermediate value between extrema. For I = 6, a symmetrical pattern emerges with a frequency of roughly 5 days.



FIG. 2.5. The dependence of the number of fast modes on system stability. Time series for x_3 as the number of fast modes J is varied for I = 6 (left panel) and I = 8 (right panel). J = 8 is the dashed line, J = 16 is the thin solid line, and J = 40 is the thick solid line.



FIG. 2.6. The dependence of the number of slow modes on system stability. Time series for x_3 as the number of slow modes I is varied. For all panels, J = 16.

System vs. Model

As mentioned above, each x_i in the system is coupled to J small-amplitude fast variables. A 10-day time series for a single x_i with its corresponding y_i^j 's is shown in Fig. 2.7 (I = 8, J = 4). In the top frame we compare the time series of x_1 with and without coupling from the fast variables. The effect from the y_1^j variables has an increasing effect on the difference between the two integrations for x_1 during the forecast. In fact, for the first six days of the forecast, there is no difference between the two. By day 10 however, the values differ by half the climatological span of x_1 . Note



FIG. 2.7. A 10 day time series for x_1 and $y_1^{1,2,3,4}$ using I = 8, J = 4. For the top frame, the solid line represents the time series for x_1 with fast mode coupling (the system). The dashed line has fast mode coupling turned off, and thus is integrating Equation (1) with h = 0. The bottom four frames are the fast mode time series coupled to x_1 .

that the fast variables all vary differently throughout the forecast, with amplitudes on the order of 10% those of the slow variables. Disturbances tend to propagate in about 4 days.

III. METHODS

From a given initial condition, the trajectory of the truth (\mathbf{z}^a) is created by integration of the system (i.e. using both the slow and fast variables in Equations (1) and (2)). A two and three dimensional view of this attractor is shown in Fig. 3.8. The forecast for each ensemble member is then completed by setting h = 0.5 in the governing equations (hereafter referred to as the 'model'). In other words, the model is rendered imperfect by dampening the effect of the fast modes in Equation (2) by 50%. The same two and three dimensional slices can be seen in Fig. 3.9, now for the model. 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. For both the truth and model forecasts, integration of the differential equations is completed using the fourth order Runge-Kutta method with a time step of 0.01. Rigorous shadowing attempts would be made using far more advanced methods of integration, with much smaller time steps. However, for the purpose of this study of short forecasts, the difference is negligible.



FIG. 3.8. System Attractor. The left panel shows a two-dimensional view of the system attractor looking at x_1 vs x_3 . The right panel shows a three-dimensional view using x_1, x_2 , and x_4 (I = 4, J = 16).



FIG. 3.9. Model attractor. The left frame shows a two-dimensional view of a typical forecast looking at x_1 vs x_3 . The right frame shows a three-dimensional view of the forecast using x_1, x_2 , and x_4 (I = 4, J = 16).

Ensemble Creation

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 from each of these states was determined using a 50-day integration of the system. The goal is then to use the model to shadow these trajectories with an ensemble of 20 members.

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

$$C^{\text{init}} = \frac{0.05^2 \sigma_{\text{clim}}^2}{\lambda} C \tag{3}$$

where λ is the average eigenvalue of C and σ_{clim} is the climatological standard deviation. First, a control state is picked in each hypersphere by adding appropriately distributed random noise to the I slow modes of the truth as follows:

$$\mathbf{z}_{0}^{f}(1:n,1) = \mathbf{z}_{0}^{a}(1:n,1) + \sqrt{C^{\text{init}}}\mathbf{y}(I,1).$$
(4)

The vector \mathbf{y} is an *I*-dimensional vector consisting of random entries from a Gaussian distribution. The Cholesky decomposition is used to calculate the square root of C^{init} [23]. The remaining 19 ensemble members are then chosen using the same method, but using the control state $\mathbf{z}_0^f(1:n,1)$ as the central reference point.

$$\mathbf{z}_{0}^{f}(1:n,j) = \mathbf{z}_{0}^{f}(1:n,1) + \sqrt{C^{\text{init}}}\mathbf{y}(I,1) \quad \text{for } j = 2, 3, \dots, 20.$$
(5)

Making Predictions

Once the ensemble of initial states has been created for each hypersphere, the trajectory of each ensemble member is forecast using the model. However, every 0.02 time units (2 time steps), the ellipsoid encompassing the ensemble members is analyzed to determine expanding and contracting directions (i.e. those directions with positive/negative local finite time Lyapunov exponents). These lengths and directions are calculated using singular value decomposition (SVD) [23]. Let $\mathbf{s} = s_1, s_2, \cdots, s_I$ be the singular values, with $U = [\mathbf{u}_1, \mathbf{u}_2, \cdots, \mathbf{u}_n]$ the matrix of left-singular vectors. At each time interval, the vector \mathbf{s} is compared to \mathbf{s} from the previous interval to determine which singular values are decreasing. Corresponding directions are matched between time steps using the dot products of the vectors composing U. A two dimensional analog is shown in Fig. 3.10. Note that all ensemble members lie within the ellipse, and the control state is located at the center. SVD allows us to identify the ellipse semi-axes $s_1\mathbf{u}_1$ and $s_2\mathbf{u}_2$.



FIG. 3.10. A two-dimensional representation of a 20-member ensemble (I = 2). Let A be the 2 × 20 matrix of ensemble members. The semi-axes are calculated using SVD, where s_1 and s_2 are the two singular values of the matrix A.

An example one-dimensional 50 day ensemble forecast is shown in Fig. 3.11 (I = 6, J = 16). Note that for the first 20 days, almost all ensemble members remain close to the truth. However, by day 30, the spread of ensemble members covers the entire state space, and the ensemble no longer accurately represents the truth.



FIG. 3.11. A 20-member ensemble forecast for x_3 initialized using Equation (5). The black line represents the truth, while the blue lines are individual ensemble members (I = 6, J = 16).

IV. 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. 4.12. The ellipse representing the ensemble members diverges from the trajectory of the truth when the system quickly turns in a new direction. Similarly, for the larger system the trajectory of the truth experiences sudden changes when contracting directions start to expand. This can be seen by plotting the number of expanding directions over a 50 day forecast for a particular hypersphere as in Fig. 4.13 (I = 6, J = 16). This unstable dimension variability is well documented in the literature [5, 6, 24, 25].



FIG. 4.12. Schematic of forecasting in two dimensions. The black line represents the trajectory of the truth (\mathbf{z}^a) . $\mathbf{z}_{T_i}^f$ is the ellipsoid encompassing the ensemble members at time T_i . Note that the time steps T_0, T_1, T_2, T_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 $\mathbf{z}_{T_i}^a$. Here forecasting fails at time T_2 .

Variance inflation is an attempt to improve the ensemble forecast by inflating the ellipsoid along contracting dimensions of the state space. For each direction i in which the ensemble is contracting, ensemble members are inflated along the i^{th} semi-minor



FIG. 4.13. The number of expanding directions over a 50 day forecast. Clearly the dynamics are changing as the state wanders through the attractor (I = 6, J = 16).

axis of the ellipsoid. Although this adds uncertainty, Danforth and Yorke argue that the amount is minimal [5]. If the ellipsoid continues to contract in these dimensions, the uncertainty introduced continuously decreases. However, if the system begins to expand along a previously contracting dimension, the ensemble has captured some of the change. This concept is illustrated for two dimensions in Fig. 4.14.



FIG. 4.14. Schematic of variance inflation in two dimensions. The black line represents the trajectory of the truth. $\mathbf{z}_{T_i}^f$ is the ellipsoid encompassing the ensemble members at time T_i . At each time step, the ellipsoid is inflated by φ along contracting directions. Note that forecasting now achieves its goal in contrast to Fig. 4.12.

As described above, the directions of the semi-axes are calculated using SVD throughout the forecast at periodic time intervals. At each point, **s** is compared to **s**', where **s**' represents the singular values for the previous time interval. For each i in which $s_i < s_i'$, each ensemble member is inflated in direction $\mathbf{b} = s_i \mathbf{U}_i$.

Let φ be the inflation amount, whose magnitude will be discussed later, and let A be the matrix of ensemble members minus the control state prior to inflation. Define

$$M = I_{m \times m} + \frac{(\varphi - 1)}{(||\mathbf{b}||_2)^2} \mathbf{b} \mathbf{b}^{\top}$$
(6)

where $I_{m \times m}$ is the *I*-dimensional identity matrix (i.e. m = I) and \mathbf{b}^{\top} is the transpose of **b**. We now define A' = MA. Adding the control state to each vector in A' then yields ensemble members inflated by φ in dimension U_i .



FIG. 4.15. Successful inflation in a forecast. The black line is the trajectory \mathbf{z}^a , the red lines are ensemble member forecasts without inflation, and the blue lines are forecasts with inflation. At point A, forecasting fails without inflation as noted by the end of the red lines. Point B represents a second inflation of the ensemble (I = 6, J = 16).

The benefit of inflation in a particular forecast has been isolated in Fig. 4.15 for dimensions x_2 and x_3 (I = 6). Note that as the trajectory of \mathbf{z}^a begins to turn, the forecasts (blue and red lines) begin to diverge from the black line. As represented by the end of the red lines, shadowing fails at point A without inflation, when there is no longer overlap between the σ -ball surrounding \mathbf{z}^a and the ensemble ellipsoid. However, the blue lines are inflated at point A in previously contracting directions. As the figure illustrates, these shifts help capture the change in direction for \mathbf{z}^a and shadowing survives. A second inflation for the blue lines occurs at Point B. Note that since there is no change in the direction of \mathbf{z}^a , at first, there is more uncertainty in the ellipsoid encompassing the ensemble members. However, as \mathbf{z}^a continues in this direction, the uncertainty is dampened.

Forecasting Results

For each of the 500 initial hyperspheres, a 50-day model forecast for the 20 ensemble members was created with and without inflation. For each hypersphere, the mean of the 20 ensemble members was compared to the truth, and the root mean square error (RMSE) and anomaly correlation (AC) were calculated as follows:

$$RMSE = \sqrt{\sum_{i=1}^{I} (\mathbf{z}^f - \mathbf{z}^a)^2}$$
(7)

$$AC = \frac{(\mathbf{z}^f - \overline{\mathbf{z}}) \cdot (\mathbf{z}^a - \overline{\mathbf{z}})}{\|\mathbf{z}^f - \overline{\mathbf{z}}\| \|\mathbf{z}^a - \overline{\mathbf{z}}\|}$$
(8)

where \overline{z} represents the vector of climatological averages for each slow variable. RMSE and AC calculations were then averaged over all hyperspheres for the duration of each forecast. Representative plots for I = 4, 5, and 6 are shown in Figs. 4.16 and 4.17 with an inflation factor of 1%. Note that for all plots, forecasts created using inflation appear to diverge more slowly from the truth than forecasts without inflation. This effect is stronger for greater I, where variance inflation lowers RMSE and increases AC. For the given system, a RMSE around 9 represents saturation, at which point there is no correlation between the forecast and the truth. For I = 4, RMSE overshoots this saturation value before returning to 9. In fact, both RMSE and AC exhibit unexpected behavior by day 20. This is most likely a result of having too few locations around the latitude band, leading solutions to go out of phase and then back in phase accidentally. The black lines in Figs. 4.16 and 4.17 indicate a large difference between the three systems. Without inflation, the model better forecasts the truth for larger I, as can be seen with the black line shifting to the right.



FIG. 4.16. Representative averaged RMSE for different values of I. The dashed line represents no inflation, while the solid line is with inflation of 1%. For all panels, J = 16.



FIG. 4.17. Representative averaged AC for different values of I. The dashed line represents no inflation, while the solid line is with inflation of 1%. The dotted horizontal line represents the 0.6 AC threshold. For all panels, J = 16.

Similarly, Figs. 4.18 and 4.19 respectively show averaged RMSE and AC plots for I = 6, J = 16, while φ is varied. The trend for these plots is more complicated to analyze. Clearly $\varphi = 0.5\%$ shows the least improvement with variance inflation as φ is too small to have a meaningful effect. The RMSE plots seem to indicate variance inflation has a greater effect for $\varphi = 2\%$, although the AC plots are less convincing. One potential reasoning for greater inflation yielding less improvement is that the inflation is in the wrong direction (i.e. away from the truth). By inflating the ellipsoid in a potential incorrect direction, variance inflation is introducing more error, thereby reducing any benefit gained.



FIG. 4.18. Representative averaged RMSE for different inflation amounts φ . The dashed line represents no inflation, while the solid line is with inflation. For all panels, I = 6, J = 16.



FIG. 4.19. Representative averaged AC for different inflation amounts φ . The dashed line represents no inflation, while the solid line is with inflation. The dotted horizontal line represents the 0.6 AC threshold. For all panels, I = 6, J = 16.

A forecast is considered to have failed once AC drops below 0.6, and yields an estimate of the shadowing time [3]. For each state space size, the averaged shadowing 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, respectively. Variance inflation was deemed to have been successful if the duration for an acceptable forecast improved by more than 5% of the average shadowing time for forecasts with $\varphi = 0$. Similarly, the technique was considered to have failed if the duration worsened by the same threshold. If there was any measurable improvement with inflation, forecasts were denoted "inflation helped." A second metric employed is a count of the number of hyperspheres that shadowed the truth for a particular time interval. For each forecast, the goal was set to 10% beyond the average shadowing time. Each experimental setup was run in triplicate, averaging the counts for all metrics. Results for I = 4, 5, and 6, with inflation amounts (φ) of 0.5%, 1%, 2%, and 5% are recorded in Table 4.1 with totals out of 500. For all experiments J = 16 to ensure the truth had somewhat of a regular pattern as described above in Fig. 2.6.

Table 4.1 indicates that variance inflation has a greater effect for larger I, with the best results exhibited for I = 6. In fact, the number of hyperspheres in which variance inflation helped was an order of magnitude better for I = 5 than for I = 4. These numbers doubled for I = 6 relative to I = 5, likely due to the increased regularity in \mathbf{z}^a for I = 6 observed in Fig. 2.6. 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, it might be evidence that error introduced by inflation in directions that continuously contract has a minimal effect as argued by Danforth and Yorke [5]. This effect would likely be strongest in systems with larger I, where the expanding dimensions have a more dominant role.

Naturally, the success of variance inflation is also strongly dependent on φ , with increasing φ generally corresponding to greater improvement. By increasing the

Dimension	φ	Inflation	Inflation	Inflation	Reached Goal	Reached Goal
I		Worked	Helped	Failed	Inflation	Non-inflation
4	0.5%	0	6	1	224	224
4	1%	1	7	1	224	224
4	2%	3	9	2	224	225
4	5%	3	16	4	225	226
5	0.5%	7	51	7	211	211
5	1%	13	84	9	214	213
5	2%	20	115	19	210	209
5	5%	34	156	37	214	210
6	0.5%	5	116	4	102	103
6	1%	86	273	27	456	456
6	2%	22	244	19	104	105
6	5%	59	295	59	117	104

TABLE 4.1. Variance inflation forecasting results out of 500 hyperspheres

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, variance inflation can make a forecast worse. By inflating in directions differing from the truth, the method introduces additional error, which can be quite significant. The number of hyperspheres in which variance inflation failed also naturally increases with φ . The more inflation away from the truth, the worse the forecast. However, as mentioned above, these directions are contracting, and thus, the error introduced is minimal.

The total number of hyperspheres for which shadowing succeeded was relatively constant between inflation and non-inflation forecasts. One possible explanation is that the effectiveness of inflation is not dependent upon the duration of the original (non-inflated) forecast. For example, inflation is equally likely to help improve (or hurt) a forecast that shadows the truth for 3 days or 20 days. Only for $I = 6, \varphi = 5\%$ is there a marked improvement in the number of hyperspheres reaching the desired shadowing time. Note that the experiment run for $I = 6, \varphi = 1\%$ appears to be an outlier from the observed trend, exhibiting the most improvement of any with a 17.2% success rate. These conditions might be the unique balance in which the positive effects of variance inflation are maximized with additional error being limited.

As shown in Table 4.1, for many of the hyperspheres, variance inflation successfully improved the duration for which a forecast is accurate. The potential for this technique can be seen by averaging RMSE and AC over the hyperspheres in which inflation was successful as shown in Figs. 4.20 for I = 6, J = 16. Both plots show improvement relative to forecasts made with no inflation. Therefore if these opportunities can be isolated as was done in Fig. 4.15, variance inflation can have great utility.



FIG. 4.20. Averaged RMSE and AC for hyperspheres in which inflation improved forecast. The black line represents no inflation, while the blue line is with inflation of 1% (I = 6, J = 16).

Monte Carlo Forecasting Results

To examine the trends for larger dimensional systems, experiments were repeated for $I = 4, 6, 8, \dots, 18$. To increase beyond I = 18, more ensemble members would be needed in order for SVD to be applicable. However, due to computation time limitations for estimating the local dynamics of the system attractor, initial ensemble members were chosen using a Monte Carlo approach. Ensemble members were chosen to be within 0.5% of the climatological span of the attractor of the initial state. The same forecasts and analyses were completed as in the smaller dimensional experiments. The average shadowing times for non-inflated forecasts are shown in Fig. 4.21. Although the times agree with the attractor-based ensemble experiments for I = 4, the Monte Carlo experiments had a significantly higher average shadowing time for I = 6 dimensional systems.



FIG. 4.21. The average shadowing time for non-inflated forecasts using the Monte Carlo approach for I = 4 to I = 18 dimensional systems.

After normalizing by the times in Fig. 4.21, the number of hyperspheres whose forecasts reached their goal with and without inflation are enumerated in Fig. 4.22. For most systems, there is little difference between the number of hyperspheres reaching the goal with and without inflation. However, I = 6 and I = 16 are notable exceptions, especially for greater inflation amounts. Too much inflation in highly regular systems can introduce too much error and forecasting fails. Since this was not observed previously (for I = 6), this result is also on account of having initial ensemble members varying too much from the initial truth, particular in directions away from the system attractor. Inflation of 2% or 5% leads to the ensemble quickly diverging from the truth. For I = 4, the counts with the Monte Carlo approach are lower than those in Table 4.1. This may indicate that the forecasts that do reach their goal with the Monte Carlo method, greatly exceed it, thereby raising the average shadowing period, and hence higher goal. For I = 6, the results are difficult to compare because of the outlier in Table 4.1.



FIG. 4.22. Number of hyperspheres reaching the desired shadowing period of 10% beyond the $\varphi = 0$ average using the Monte Carlo approach. Solid lines represent forecasts with no inflation, while dashed lines are with variable inflation.

Counting the number of hyperspheres for which inflation worked and failed with the Monte Carlo technique, the trends observed in Table 4.1 do not tell the full story as seen in Fig. 4.23. In fact, I = 6 and I = 16 dimensional systems are unique in that small amounts of inflation can greatly improve forecasts. However, for these same systems, too much inflation can lead to an increase in the number of failures. It is in these regular systems, as discussed previously, for which the model is most sensitive to variance inflation. Note that in agreement with Table 4.1, the best improvement occurs with 1% inflation on a 6-dimensional system. Between these two values for I, the trend is reversed for the number of improvements. More inflation leads to an increase in the number of hyperspheres for which inflation works. As expected, this also corresponds with relatively fewer failures, although the number of times inflation works and fails is quite similar for these systems. For all systems, increasing the inflation fails.



FIG. 4.23. Number of hyperspheres where inflation worked and failed using the Monte Carlo approach on I = 4 to I = 18 dimensional systems. Panel (a) shows the number of hyperspheres for which inflation improved the forecast by over 5% of the average shadowing time. Panel (b) shows the number for which inflation worsened the forecast by the same amount.

V. SHADOWING AND STALKING

The criteria of shadowing is a different perspective for assessing the quality of a model's predictions. The trajectory of the truth is known a priori to within a given uncertainty (σ), and whose location at a given time is represented by an Idimensional ball of radius σ . At regular intervals, the intersection between the ellipsoid encompassing the ensemble members and the σ ball is calculated. The ensemble ellipsoid is then redefined to approximate the intersection, while the other ensemble members are ignored. Provided this intersection is nonempty, the trajectory of some ensemble members yield accurate representations of reality. A schematic of this technique is shown in Fig. 5.24. 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.

Stalking Results

The same initial conditions utilized in the forecasting experiments were used to create 50-day predictions. However, every 0.04 time units (every 4th time step), the ensemble ellipsoid was redefined as described above. This interval was chosen to ensure contracting dimensions could still be calculated without the redefinition altering semi-axes directions beyond recognition. A new collection of 20 ensemble members were then chosen lying within the redefined ellipsoid ($\sigma = 10\%$ of the climatological range). As before, RMSE and AC were calculated and compared for both shadowing and stalking experiments. If the time period for which AC remained



FIG. 5.24. Schematic of the redefining of an ensemble in two dimensions. The star represents the known location of the truth. The solid squares are supplemented by a collection of points meant to approximate the overlap of the ellipse and the ball around the truth. The open circles are previous ensemble members that will now be ignored.

above 0.6 improved by more than 5% of the average shadowing time, then stalking was considered successful. Further, the number of hyperspheres σ -shadowed/ σ -stalked for at least 10% beyond the average shadowing time were counted. Results for I = 4, 5, and 6, with inflation amounts (φ) of 0.5%, 1%, 2%, and 5% are given in Table 5.2.

As with forecasting, the shadowing experiments exhibited a greater benefit from inflation with larger dimensions, although the difference is less than before. However, the number of hyperspheres for which inflation worked does not seem to depend upon φ over the range 0.5% - 5%. Since the number of hyperspheres in which inflation worked 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 worked is nearly equivalent to the number for which inflation failed for all experiments. Thus, inflation can also have a negative effect upon shadowing.

Dimension	φ	Inflation	Inflation	Inflation	Reached Goal	Reached Goal
I		Worked	Helped	Failed	Stalking	Shadowing
4	0.5%	57	65	55	376	375
4	1%	52	58	54	375	378
4	2%	60	69	52	376	368
4	5%	56	65	56	374	374
5	0.5%	63	74	59	330	326
5	1%	66	73	64	337	335
5	2%	65	73	57	335	328
5	5%	72	80	56	337	320
6	0.5%	55	68	60	233	238
6	1%	73	84	62	328	317
6	2%	64	83	61	231	227
6	5%	71	83	51	260	240

TABLE 5.2. Shadowing and stalking results out of 500 hyperspheres

The number of 50-day σ -shadowed and σ -stalked hyperspheres tells an interesting yet complicated story. Examining the non-inflation results alone indicates that shadowing is most successful for smaller dimensional systems. With fewer expanding dimensions, redefining the ensemble ellipsoid every 0.04 time units can readily ensure some ensemble members remain close to the truth for the entire 50 day experiment. Stalking results demonstrate that inflation has mixed success, although the overall effect is quite mild. For many experiments, the change was less than five hyperspheres in either direction. The most improvement occurred with maximal inflation ($\varphi = 5\%$), with increases of 17 and 20 hyperspheres for I = 5 and I = 6, respectively. A notable outlier occurs for I = 6, $\varphi = 1\%$ where the number of hyperspheres both σ -shadowed and σ -stalked is noticeably higher than other experiments. This might be a due to the regularity for I = 6 observed in \mathbf{z}^a , discussed previously. Under these ideal initial conditions, stalking improves the number of accurate predictions by 11 hyperspheres. These results indicate the stalking can be an improvement over shadowing, however, inflation in directions away from the truth can be harmful.



FIG. 5.25. Number of hyperspheres reaching the desired shadowing period of 10% beyond average using the Monte Carlo approach. Solid lines represent forecasts with no inflation, while dashed lines are with variable inflation.

As with forecasting, stalking experiments were extended to greater dimensional systems $(I = 4, 6, 8, \dots, 18)$ using a Monte Carlo approach, where initial ensemble members were chosen without knowledge of the local dynamics of the system attractor. The same redefining technique and analysis described above was applied to these systems. Fig. 5.25 depicts the number of hyperspheres shadowed and stalked for at least 10% beyond the average shadowing time. Many of these hyperspheres even remained close to the truth for the entire 50-day experiment. Unlike the results presented in Table 5.2, it appears as though the number of hyperspheres shadowing successfully increase from I = 4 to I = 6. However, these two data points are significantly higher than for all other dimensional systems. This may indicate that many of the hyperspheres would have remained close to the truth for all experimental setups, there is hardly a difference between stalking and shadowing (i.e. inflated and non-inflated) experiments. The only variation is with $\varphi = 5\%$ for I = 4 and I = 6, in agreement with Table 5.2.



FIG. 5.26. Number of hyperspheres where inflation worked and failed using the Monte Carlo approach on I = 4 to I = 18 dimensional systems. Panel (a) shows the number of hyperspheres for which inflation improved the forecast by over 5% of the average shadowing time. Panel (b) shows the number for which inflation worsened the forecast by the same amount.

The number of hyperspheres for which inflation worked and failed using the Monte Carlo approach for I = 4 - 18 is shown in Fig. 5.26. The number for which inflation worked is slightly greater than the number for which inflation failed, although the counts are close for all experiments. For I = 4 and I = 6, both counts are significantly higher than observed in Table 5.2. Interestingly, the trends detected for forecasting results are quite different. In fact, greater inflation generally corresponds to lower number of hyperspheres for which inflation failed. This might be signifying dot product errors in correlating ensemble semi-axes directions between time steps for lower inflation amounts. With greater inflation, the model can better determine corresponding directions during redefinition of the ensemble. The number of hyperspheres for which inflation worked is relatively constant for all inflation amounts. Note that the extremal values for I = 6 and I = 16 for forecasting experiments are not as distinguishable. By redefining the ensemble every 0.04 time steps, the large scale regularity for those systems is important. Since the truth is known throughout, ensemble members are forced to remain nearby regardless of the shape of the attractor.

VI. CONCLUSIONS

Using the Lorenz '96 coupled system as an analog for atmospheric dynamics, we were able to analyze the potential for using inflation in models 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 at periodic time steps. In the former the inflation technique is referred to as variance inflation, and in the later, stalking. To expand the results for larger dimensional systems, a Monte Carlo approach was used to simplify the definition of the initial ensemble. Under idealized conditions, inflation was shown to be beneficial for both techniques.

Overall, the greatest determinant of a forecast's success was the tuning of the model's parameters. For the Lorenz '96 model, slightly adjusting the number of slow and fast modes, and the degree of coupling between them, one could vary between regular and highly chaotic systems. This system had the greatest regularity for I = 6 and I = 16. However, by increasing the number of fast modes (J), greater regularity can be achieved for a fixed number of slow modes. In contrast to hyperbolic systems, the present system exhibits unstable dimension variability and fluctuating Lyapunov exponents as is common in higher dimensional systems [25].

For forecasting experiments, inflation showed the potential to work as illustrated in Fig. 4.15. Inflation in directions away from the truth, however, can worsen a forecast. As a result, greater inflation amounts always corresponded to an increase in the number of hyperspheres for which inflation failed. For regular systems such as I = 6 and I = 16, smaller inflation amounts can often successfully improve a forecast, while the error introduced from larger inflation amounts is too great. However, for less regular systems with dimensions between 6 and 16, 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 trajectory that remains close to the truth for a given time and assess how inflation effects this trajectory. For these experiments, results were relatively constant for systems with greater than 6 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.

When shadowing physical systems such as the atmosphere, one is presented with the challenges of uncertainty in the initial condition, sensitive dependence upon these initial conditions, and model error. Model error is even estimated to dominate the forecast error for the first three days in weather systems [20]. 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. However, isolating when inflation will help and when it will be harmful has yet to be established.

APPENDIX A: UNDERSTANDING SVD

Consider the $I \times (m-1)$ matrix A, where *m* represents the number of ensemble members. The *i*th column of A represents the difference between the *i*th ensemble member and the control state. Let $A = USV^{\top}$ be a singular value decomposition of A, where s_i form the diagonal entries of *S*. The eigenvalues of (AA^{\top}) equal S^2 , and increase with the number of ensemble members *m*. Although, **s** grows proportionally with \sqrt{m} , the magnitude of s_i is irrelevant for the purposes of inflation. We are solely interested in the dynamics of the semi-axes directions. (i.e. Is s_i contracting relative to the previous time interval?). Note that the magnitude of **s** represents the range of possible ensemble members, which does increase with *m*.

Consider the matrix A acting on the *m*-dimensional unit sphere. For m = I, the resulting matrix represents the boundaries of an ellipsoid. Now consider the case where m > I. We want to show that for an arbitrary vector \mathbf{y} existing in/on the ellipsoid, we can solve $A\mathbf{x} = \mathbf{y}$. We know that the size of the singular values of the ellipsoid (and thus the size of the ellipsoid) will be growing with \sqrt{m} . For simplicity, let I = 2, and the *m*-dimensional unit vectors are mapped into 2 dimensions by A. Intuitively, as *m* increases, the column space of A should be larger. Let $\mathbf{y} = \{y_1, y_2\}$, $\mathbf{x} = \{x_1, x_2, \ldots, x_m\}$, and A_{ij} be the entry (i, j) of A. Then we are trying to solve

$$A_{11}x_1 + A_{12}x_2 + \dots + A_{1m}x_m = y_1 \tag{9}$$

$$A_{21}x_1 + A_{22}x_2 + \dots + A_{2m}x_m = y_2 \tag{10}$$

Ensuring x is a unit vector can be accomplished via scaling. Clearly, given \mathbf{y} , we can find infinitely many vectors \mathbf{x} such that $A\mathbf{x} = \mathbf{y}$.

APPENDIX B: MATLAB CODE

All experiments were computed using the University of Vermont Bluemoon cluster (an IBM e1350 High Performance Computing system) run by the Vermont Advanced Computing Center. The cluster is made possible by grants from the National Aeronautics and Space Administration (NASA) with strong support from U.S. Senator Patrick Leahy and Vermont EPSCoR. The original matlab code can be found at: http://www.uvm.edu/ cdanfort/nolink/atmos/stalking-code.tar. It was used to run 264 experiments comprising of 11,458 submitted jobs.

LITERATURE CITED

- E. N. Lorenz, "A study of the predictability of a 28-variable atmospheric model," Tellus, 17, 321–333 (1965).
- [2] C. E. Leith, "Objective methods for weather prediction," Ann. Rev. Fluid Mech., 10, 107–128 (1978).
- [3] E. Kalnay, Atmospheric Modeling, Data Assimilation and Predictability (Cambridge Univ. Press, 2003).
- [4] Z. Toth and E. Kalnay, "Ensemble forecasting at nmc: the generation of perturbations," Bull. Amer. Meteor. Soc., 74, 2317–2330 (1993).
- [5] C. M. Danforth and J. A. Yorke, "Making forecasts for chaotic physical processes," Phys. Rev. Letters, 96, 144102 (2006).
- [6] S. Dawson, C. Gregobi, T. Sauer, and J. Yorke, "Obstructions to shadowing when a lyapunov exponent fluctuates about zero," Phys. Rev. Letters, 73, 1927–1930 (1994).
- [7] D. V. Anosov, "Geodesic flows and closed riemannian manifolds with negative curvature," Math, 90 (1967), proc. Steklov Inst.
- [8] R. Bowen, "ω-limit sets for axiom a diffeomorphisms," J. Differential Equations, 18, 333–339 (1975).
- C. Robinson, "Stability theorems and hyperbolicity in dynamical systems," Rocky Mountain J. of Math., 7, 425–437 (1977).
- [10] K. R. Meyer and G. R. Sell, "An analytic proof of the shadowing lemma," Funkcialiaj Ekvacioj, **30**, 127–133 (1987).
- [11] T. Sauer and J. A. Yorke, "Rigorous verification of trajectories for the computer simulation of dynamical systems," Nonlinearity, 4, 961–979 (1991).
- [12] B. A. Coomes, H. Kocak, and K. Palmer, "A shadowing theorem for ordinary differential equations," Z. Angew. Math Phys., 46, 85–106 (1995).

- [13] C. Grebogi, S. M. Hammel, J. A. Yorke, and T. Sauer, "Shadowing of physical trajectories in chaotic dynamics: Containment and refinement," Phys. Rev. Letters, 65, 1527–1530 (1990).
- [14] S. M. Hammel, J. A. Yorke, and C. Grebogi, "Do numerical orbits of chaotic dynamical processes represent true orbits?" J. of Complexity, 3, 136–145 (1987).
- [15] S. M. Hammel, J. A. Yorke, and C. Grebogi, "Numerical orbits of chaotic processes represent true orbits," Bull. Am. Math. Soc., 19, 465–469 (1988).
- [16] E. N. Lorenz, "Predictability: A problem partly solved," (1996), proc. ECMWF Seminar on Predictability.
- [17] D. S. Wilks, "Effects of stochastic parametrizations in the lorenz '96 system," Q. J. R Meteorol. Soc., 131, 389–407 (2005).
- [18] D. Orrel, "Role of the metric in forecast error growth: How chaotic is the weather," Tellus, 54A, 350–362 (2002).
- [19] C. M. Danforth and E. Kalnay, "Using singular value decomposition to paramterize state-dependent model errors," J. Atmos. Sci., 65, 1467–1478 (2008).
- [20] D. Orrel, L. Smith, J. Barkmeijer, and T. N. Palmer, "Model error in weather forecasting," Nonlin. Proc. Geophys., 8, 357–371 (2001).
- [21] E. N. Lorenz and K. A. Emanuel, "Optimal sites for supplementary weather observations: Simulation with a small model," J. Atmos. Sci., 55, 399–414 (1998).
- [22] A. J. Simmons, R. Mureau, and T. Petroliagis, "Error growth estimates of predictability from the ecmwf forecasting system," Quart. J. Roy. Meteor. Soc., 121 (1995).
- [23] G. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. (The John Hopkins University Press, 1996).
- [24] G.-C. Yuan and J. Yorke, "An open set of maps for which every point is absolutely nonshadowable," Proceedings of the American Mathematical Society, 128, 909 (2000).
- [25] T. Sauer, "Shadowing breakdown and large erros in dynamical simulations of physical systems," Phys. Rev. E, 65, 36220 (2002).