Sequential Data Assimilation by Nudging

Marcel Oliver

October 2004

Abstract

Data assimilation is the process of initializing a forecast (for example a weather forecast) from incomplete observation. We consider a system described by system of differential equations where not all degrees of freedom are observable, either because this is not practical, as, for example, in the atmosphere where wind speed and temperature information above ground level is expensive to measure, or because these degrees of freedom are fundamentally inaccessible. We also assume that the system of differential equations is an accurate model of the real physics (leaving aside issues of model uncertainty), and that measured quantities are exact (leaving aside measurement errors). Thus, the task is to use the differential equation to determine the complete state at a fixed time from time series data on the accessible degrees of freedom.

The simplest possible form of this so-called *sequential data assimilation* is called *nudging* or *dynamical relaxation*—we introduce additional forces into the differential equation that "nudge" the solution toward the correct internal state as time evolves. These lectures will explore nudging in one of the simplest possible setting for which the technique is interesting: the three-dimensional Lorenz equations with one accessible observable.

Contents

Instructions for students: The homework exercises are marked "Homework" and are due one week after the last lecture of this module. The questions titled "Project" are suggestions for the term paper, or general food for thought.

1 A crash course in differential equations

1.1 Basic setting and examples

Let $D \subset \mathbb{R}^n$ be the set of admissible states, which we call the *phase space*. For a given *vector field* $f \colon D \to \mathbb{R}^n$ and *initial value* $y_0 \in D$, we seek a continuously differentiable function $y \colon [0,T) \to D$ which satisfies

$$\dot{y}(t) = f(y(t)) \,, \tag{1.1a}$$

$$y(0) = y_0 \,. \tag{1.1b}$$

The dot shall denote derivatives with respect to t, which is common practice whenever the physical meaning of t is time. The interval of existence of a solution may be unbounded, i.e. $T = \infty$, or the solution solution may cease to exist after a finite interval of time; see the remarks on blowup in Section 1.2 below. We will often drop explicit arguments and write $y = y(t)$. On the other hand, when the dependence on the initial value is important, we sometimes write $y(t; y_0)$.

Remark 1. The form of equation (1.1) seems special in several respects. However, the formulation is actually very general, as many more differential equations can be reduced to this particular case. For example:

(i) Our vector field f does not explicitly depend on t. Such equations are called *autonomous*. Equations with explicit t-dependence are called *non-autonomous*. By regarding t as the $(n+1)$st component of the unknown function which satisfies the equation $\dot{t} = 1$, a non-autonomous equation can always be written in autonomous form. However, non-autonomous equations often warrant special treatment regardless of this formal trick.

(ii) Our system is first order—no higher order derivatives appear. Any higher order system can be written as a first order system by introducing separate dependent variables for the intermediate derivatives. This procedure is illustrated in Example 1.3 below.

(iii) The system is linear and non-degenerate in $\dot{y}$. If it is not, one can sometimes use the implicit function theorem to locally write the equation as (1.1). In general, this is difficult, and further complications arise if this process degenerates. Such issues will not be considered here.

(iv) All data is specified at the initial time; the equation is an *initial value problem*. There are also *boundary value problems* where data is specified at different points in time. Such equations often lead to quite different mathematics and will not be considered here.

Example 1. (Exponential growth or decay.) Take $D = \mathbb{R}$;

$$\dot{y} = a\, y\,, \tag{1.2a}$$
$$y(0) = y_0\,, \tag{1.2b}$$

where $a \in \mathbb{R}$ is a given number. The solution is known from calculus and can be checked by differentiation,

$$y(t) = \mathrm{e}^{at}\, y_0\,. \tag{1.3}$$

Example 2. (Linear autonomous system.) Take $D = \mathbb{R}^n$;

$$\dot{y} = A\, y\,, \tag{1.4a}$$
$$y(0) = y_0\,, \tag{1.4b}$$

where A is a given $n \times n$ matrix. We can formally write

$$y(t) = \mathrm{e}^{At}\, y_0\,. \tag{1.5}$$

It is not difficult to make sense of this expression by defining the exponential of a matrix through its power series. This will be studied in more detail in a third year lecture course on differential equations or dynamical systems. If the diagonalization or the Jordan normal form of A is known, its matrix exponential can be computed explicitly. In this sense, linear autonomous systems are "boring" as the solution procedure reduces to a standard problem of linear algebra.

Points $\bar{y} \in D$ where the solution $y(t; \bar{y})$ does not change in time are called *equilibrium points* (sometimes also *stationary points* or *fixed points* or *critical points*). We observe that if y does not change, then $\dot{y}$ must be zero. Thus, according to (1.1a), equilibrium points are characterized by $f(\bar{y}) = 0$.

Example 3. (Logistic equation.) Take $D = \mathbb{R}$;

$$\dot{y} = y - a\, y^2\,, \tag{1.6a}$$
$$y(0) = y_0\,, \tag{1.6b}$$

Solving this equation is a standard exercise in calculus. However, much can be said about the behavior of the solution without actually finding $y(t)$ explicitly.

The equilibrium points are clearly $\bar{y}_1 = 0$ and $\bar{y}_2 = 1/a$. If $0 < y < 1/a$ then $f(y) > 0$ and therefore $y(t)$ must be increasing. Outside the interval marked by the two equilibria, $y(t)$ must be decreasing. This completely characterizes the possible qualitative behavior of the solutions, as is illustrated in Figure 1.

1.2 Existence and uniqueness; blowup

Differential equations of the form (1.1) can be shown to possess unique solutions under relatively mild assumptions on the vector field f. This is the content of the following theorem.

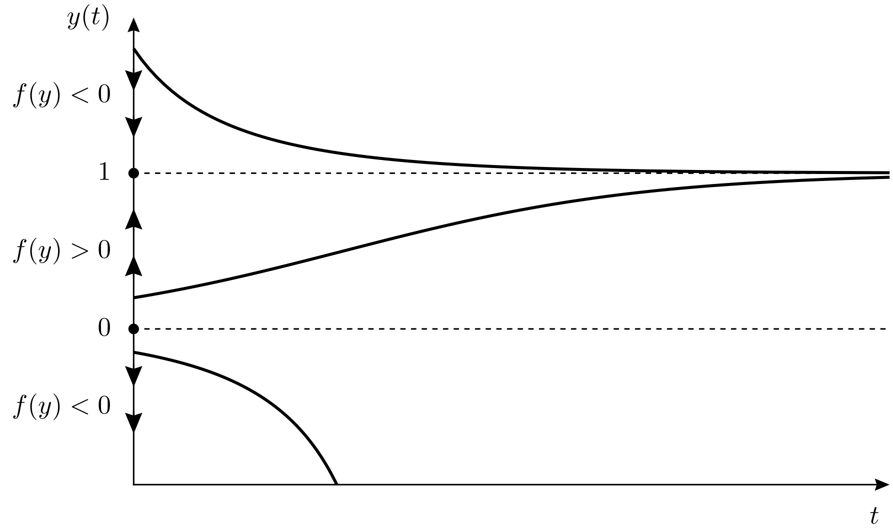

Figure 1: Solutions to the logistic equation with $a = 1$ and for three different initial values. The y-axis is the one-dimensional phase space where the direction of the vector field f is indicated. In one dimension, $y(t)$ is necessarily monotonic and the equilibrium points at $y = 0$ and $y = 1$ act as topological barriers.

Theorem 1 (Picard–Lindelöf)**.** *Assume that f is uniformly Lipshitz continuous on D, i.e. there exists some $L > 0$ such that*

$$\|f(x) - f(y)\| < L\,\|x - y\| \tag{1.7}$$

for every $x, y \in D$. Then there exists a time $T > 0$ depending only on L and $\|y_0\|$ such that

- *(i) there exists a continuously differentiable $y\colon [0,T) \to D$ which solves the ordinary differential equation* (1.1)*;*
- *(ii) the solution is unique;*
- *(iii) $y(t; y_0)$ depends continuously on y_0 for every fixed $t \in [0,T)$.*

Here and in the following we use $\|v\|$ to denote the Euclidean length of a vector $v \in \mathbb{R}^n$, i.e.

$$\|v\|^2 = v_1^2 + \cdots + v_n^2 \,. \tag{1.8}$$

The proof of this theorem is based on a fixed point argument and is thus nonconstructive. Moreover, the time interval of existence $[0, T)$ that arises in the proof of the theorem is typically much shorter than the true maximal interval of existence. However, a careful analysis of the proof (not given here) reveals that the time of existence is uniform in $\|y_0\|$. In other words, so long as the solution remains in a fixed bounded set, one can repeatedly apply the argument, thereby extending the interval of existence. The only way that the solution can cease to exist is by "blowing up", as stated more formally in the following.

Theorem 2. *Provided that f is defined on $\mathbb{R}^n$ and uniformly Lipshitz continuous on bounded subsets of $\mathbb{R}^n$, then either $T = \infty$ or*

$$\|y(t)\| \to \infty \quad \text{as } t \to T\,. \tag{1.9}$$

Example 4. Consider the differential equation

$$\dot{y} = y^2\,, \tag{1.10a}$$
$$y(0) = 1\,, \tag{1.10b}$$

Formal separation of variables gives

$$\frac{\mathrm{d}y}{y^2} = \mathrm{d}t \tag{1.11}$$

(such use of differentials can be made rigorous, or checked by direct computation on the final answer), and upon integration

$$\int_{y(0)}^{y(t)} \frac{\mathrm{d}y}{y^2} = \int_0^t \mathrm{d}t\,. \tag{1.12}$$

(The necessity to integrate from $y(0)$ to $y(t)$ on the left may seem odd, but can be derived as follows: write the left hand side of (1.11) initially as $y^{-2}\,\dot{y}\,\mathrm{d}t$, integrate from 0 to t on both sides, and then change variables from t to y on the left.) From (1.12), we obtain

$$-\frac{1}{y}\bigg|_{y(0)}^{y(t)} = t\,. \tag{1.13}$$

Solving for $y(t)$ and inserting the initial condition, we arrive at

$$y(t) = \frac{1}{1-t} \tag{1.14}$$

which clearly has blow-up time $T = 1$.

Example 5. Linear differential equations do not blow up, as solution formula (1.5) shows.

In general, we have to prove that equations do not blow up, as (most of the time) this means that the model is running out of the physically relevant regime. This has to be done for each type of equation separately. Theorems 1 and 2 suggest how to go about it: First, use the existence argument to be sure to have a solution to work with for at least a short interval of time, and then try to show that this solution remains in a bounded subset of $\mathbb{R}^n$, contradicting blowup. The latter usually requires some tricks, and we'll look at the issue in detail in Section 1.4.

Homework 1. Perform a computation similar to Example 4 that shows that the solution to the equation

$$\dot{y} = \sqrt{y}\,, \tag{1.15a}$$

$$y(0) = 0\,, \tag{1.15b}$$

is not unique. (Note that the function $f(y) = \sqrt{y}$ does not satisfy a Lipshitz condition at $y = 0$, thus Theorem 1 does not apply.)

Project 1. Write up the proof of the Picard–Lindelöf theorem, which can be found in most textbooks on differential equations. Is the requirement of the Lipshitz condition "sharp", or can it be weakened? Explore examples of non-uniqueness like Homework 1, and try to close the gap between examples and theorem.

1.3 Nonlinear equations in two dimensions

To sharpen our intuition, we will consider an equation that is only slightly more complicated than the equations discussed before. It is called *van der Pol equation*, and describes oscillations in an an electrical circuit involving a cathode ray tube. Usually written as the scalar second order equation

$$\ddot{x} + x = \mu\,(1 - x^2)\,\dot{x}\,, \tag{1.16}$$

with μ a positive parameter, it can also be written as a two-component system of first order equations by defining $y = \dot{x}$, so that

$$\dot{x} = y\,, \tag{1.17a}$$

$$\dot{y} = -x + \mu\,(1 - x^2)\,y\,, \tag{1.17b}$$

augmented with initial conditions on x and y.

The van der Pol system illustrates that even in a two dimensional phase space there are strong topological obstructions on the dynamics: For large times, an orbit can only converge to an equilibrium point, a limit cycle, or diverge. This result, called the *Poincaré–Bendixon Theorem*, should be intuitively clear; see, for example, Figure 2. A formal proof, however, is more involved and not the subject of these lectures.

1.4 The Lorenz equations and chaos

Beyond the topological constraints of dimensions one and two, very complicated behavior is possible. In the following, we will study the *Lorenz equations*

$$\dot{x} = \sigma\,(y - x)\,, \tag{1.18a}$$

$$\dot{y} = \rho x - y - xz\,, \tag{1.18b}$$

$$\dot{z} = xy - \beta z\,, \tag{1.18c}$$

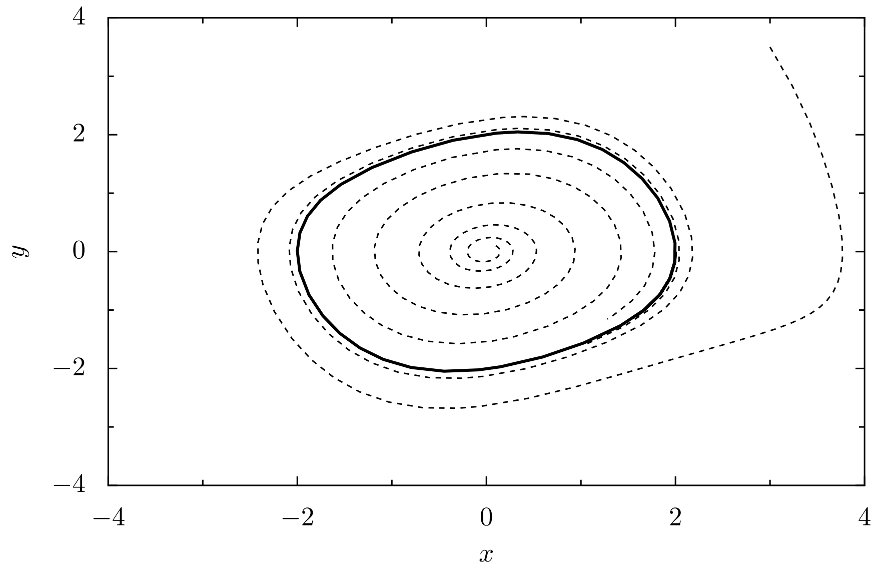

Figure 2: Phase plane of the van der Pol oscillator. The thick orbit is the stable limit cycle, the dashed lines are two other orbits that converge to the limit cycle from the inside and the outside, respectively.

where σ, ρ, and β are positive parameters. The equations represent the dynamics of the fundamental modes of a two-dimensional model of fluid convection; see [4], for example.

Part of a single trajectory in phase space is shown in Figure 3. This type of qualitative behavior is called *deterministic chaos*. The word "deterministic" indicates that each trajectory follows a well posed initial value problem and is therefore at least in principle computable. "Chaos" refers to the following.

(i) The state of the system after some time *depends sensitively on the initial data* to an extent that makes long term prediction practically impossible.

(ii) Trajectories approach a *strange attractor* as $t \to \infty$. This is a so-called fractal set with a dimension typically non-integer and larger than one. We will not go into the mathematics of these concepts here.

(iii) *Intermittency*—solutions have long relatively regular states, interrupted by intermittent periods of rapid change; see Figure 4.

Remark 2. The historic significance of the Lorenz equations is summarized in the following quote from [5].

> The significance of these equations, which were discovered by Edward Lorenz back in the 60s, is that relatively simple systems such as these could exhibit rather complex (specifically, chaotic) behavior. The chaotic

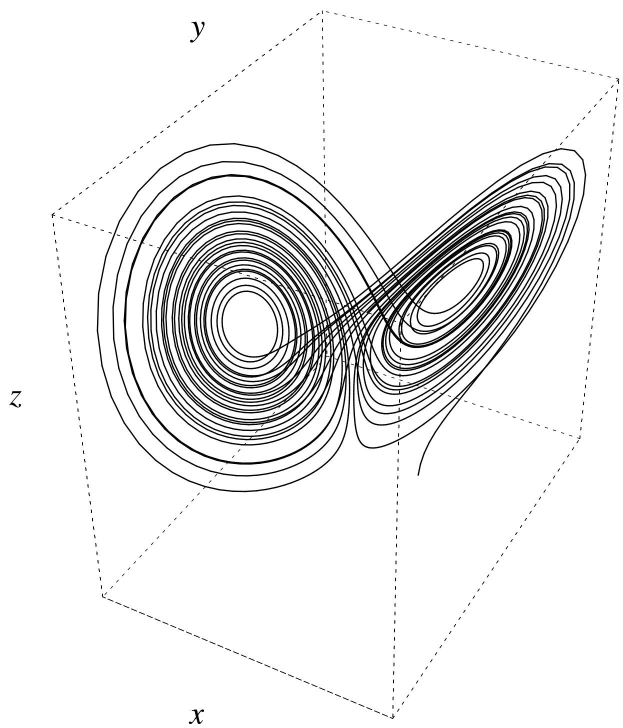

Figure 3: The Lorenz attractor. Solution curve of the Lorenz equation in phase space with $\sigma = 10$, $\rho = 35$, $\beta = \frac{8}{3}$, $x(0) = 10$, $y(0) = 0$, and $z(0) = 10$ for the time interval $t \in [0, 30]$.

aspect of this system demonstrates that, despite being given the initial conditions to any arbitrary degree of accuracy, one cannot predict a sufficiently advanced state of the system. In other words, the system has built into itself the property of amplifying small perturbations until they become so significant they affect the accuracy of the results.

For some years after Lorenz published his paper, mathematicians and physicists did not believe such behavior was even possible at all—and even for those who thought otherwise, conventional wisdom suggested that an extremely complex model would be required.

Lorenz, a meteorologist, made his discovery by observing weather phenomena—in particular, convection of fluids (and to a weatherman, the "fluid" he's most interested in is air). He took various mathematical models of fluid convection and simplified them into a system of differential equations that basically led him to the now-famous Lorenz equations.

This stripped-down system doesn't really model any real-world situations to a degree that would interest an applied scientist, but Lorenz wasn't

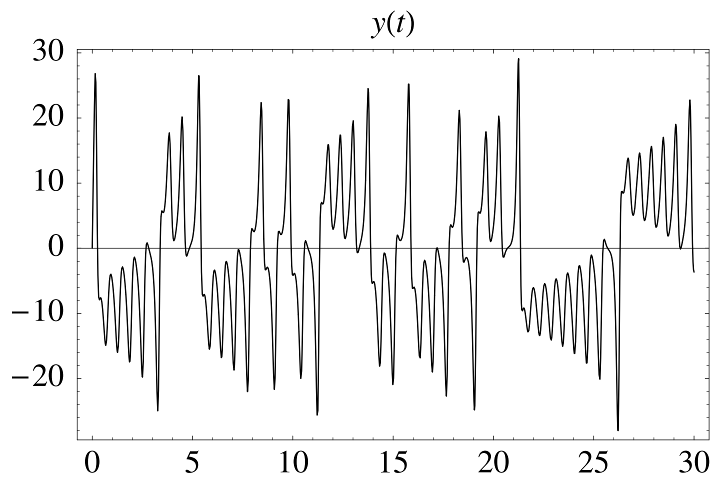

Figure 4: Time series for $y(t)$ for the solution to the Lorenz equations corresponding to the computation of Figure 3.

> concerned with modeling; rather, he wanted to examine what he felt were the fundamental, intrinsic properties of such nonlinear systems. He was more interested in the complexity of behavior than in how this behavior might translate to real-world phenomena.
>
> That said, it took years for his work to receive the attention it so well deserved. Many colleagues and peers initially either (1) refused to consider his work because, being interested in how mathematics models the world around them, they gave no thought to abstract models that "don't behave anything like real weather"; or (2) did not believe such behavior was indeed exhibited by these equations, and attributed the results to computational errors.
>
> Well, nowadays every mathematician and scientist understands and accepts the idea first put forth by Lorenz, that chaotic behavior can arise naturally from simple nonlinear systems, and by extension, real-world phenomena that are modeled by similar nonlinear systems.

In the remainder of this section, we will prove that solutions to the Lorenz equations do not blow up. In fact, we will show something stronger, namely that solutions to the Lorenz equations enter a set which is bounded independent of the magnitude of the initial data after a finite transient time (which can well depend on the initial condition).

Writing $Y = (x, y, z)$, we first note that we need to find a bound on the quantity $\|Y\|^2 = x^2 + y^2 + z^2$. The plan is to find a (scalar!) differential equation for this quantity. However, even for simple examples it is usually impossible to achieve exact closure; usually we are forced apply some inequalities to close this equation. In this case, we will eventually even replace the norm by a different dynamic quantity to obtain the strongest possible result.

We begin by multiplying the components of the Lorenz system (1.18) by x, y, and z, respectively, thereby obtaining

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}x^2 = \sigma\, xy - \sigma\, x^2\,, \tag{1.19a}$$

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}y^2 = \rho\, xy - y^2 - xyz\,, \tag{1.19b}$$

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}z^2 = xyz - \beta\, z^2\,. \tag{1.19c}$$

Adding the three equations, we obtain

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\|Y\|^2 = -\sigma\, x^2 - y^2 - \beta\, z^2 + (\sigma+\rho)\, xy\,. \tag{1.20}$$

To proceed, note the following.

Lemma 3. *For arbitrary real numbers a and b,*

$$|ab| \le \tfrac{1}{2}\, a^2 + \tfrac{1}{2}\, b^2\,. \tag{1.21}$$

Proof. The binomial formula gives

$$0 \le (a \pm b)^2 = a^2 \pm 2ab + b^2\,. \tag{1.22}$$

Solving for $\mp ab$ yields (1.21). □

We use this lemma with $a = (\sigma+\rho)\, x$ and $b = y$ to estimate the cross term on the right of (1.20),

$$|(\sigma+\rho)\, xy| \le \tfrac{1}{2}(\sigma+\rho)^2\, x^2 + \tfrac{1}{2}\, y^2 \tag{1.23}$$

and therefore

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\|Y\|^2 \le \bigl(\tfrac{1}{2}(\sigma+\rho)^2 - \sigma\bigr)\, x^2 - \tfrac{1}{2}\, y^2 - \beta\, z^2\,. \tag{1.24}$$

Depending on the values of the parameters, two cases are possible. If

$$\tfrac{1}{2}(\sigma+\rho)^2 - \sigma < 0\,, \tag{1.25}$$

then there exists a constant $\gamma > 0$ (take the smallest in absolute value of the coefficients on the right side of equation 1.24) such that

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\|Y\|^2 \le -\gamma\,\|Y\|^2\,. \tag{1.26}$$

This differential inequality has the same simple form as the differential equation in Example 1, and therefore yields an exponentially decreasing bound on the norm of Y. In other words, no matter where in the phase space we start out, the solution will converge to zero. This is obviously not a very interesting choice of parameters for the Lorenz model.

In the general case, we can still find a constant $c > 0$ (take the largest of the coefficients on the right side of equation 1.24) such that

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\|Y\|^2 \leq c\,\|Y\|^2\,. \tag{1.27}$$

This differential inequality implies an exponentially increasing bound on the norm of Y. In other words, the norm of Y can grow, but not so fast as to blow up in finite time.

However, the computer generated Figures 3 and 4 suggest that not only do trajectories not blow up, but that also they are restricted to a bounded region in phase space. Looking at the above argument once again, we see that the troublesome term is the cross-term proportional to xy on the right side of (1.20). On the other hand, this same term appears on the right hand side of the z-equation of the Lorenz model as well, and we can add an appropriate lower order term proportional to z to the norm which we intend to estimate. More precisely, we set

$$V = \tfrac{1}{2}\rho\,x^2 + \tfrac{1}{2}\sigma\,y^2 + \tfrac{1}{2}\sigma\,z^2 - 2\rho\sigma z\,. \tag{1.28}$$

The choice of prefactors in front of the quadratic terms is just a convenience, however it is crucial to find the matching coefficient in front of z.

By adding the appropriate multiples of (1.19a–c) and (1.18c), we directly obtain

$$\begin{aligned}\frac{\mathrm{d}V}{\mathrm{d}t} &= -\rho\sigma\,x^2 - \sigma\,y^2 - \beta\sigma\,z^2 + 2\beta\rho\sigma z\,.\\ &= -\big(\rho\sigma\,x^2 + \sigma\,y^2 + \tfrac{1}{2}\beta\sigma\,z^2\big) - \tfrac{1}{2}\beta\sigma\,z^2 + 2\beta\rho\sigma z\\ &\leq -\gamma\,V - \tfrac{1}{2}\beta\sigma\,z^2 + 2\rho\sigma\,(\gamma+\beta)\,z\,,\end{aligned} \tag{1.29}$$

where γ is chosen small enough such that the last inequality holds, for example $\gamma = \min\{2\sigma, 2, \beta\}$. To find a usable bound for the last term on the right of (1.29), we apply Lemma 3 with $a = 2\rho(\gamma+\beta)\sqrt{\sigma/\beta}$ and $b = \sqrt{\beta\sigma}\,z$. The resulting differential inequality reads

$$\frac{\mathrm{d}V}{\mathrm{d}t} \leq c - \gamma\,V \tag{1.30}$$

with

$$c = \frac{2\sigma\rho^2(\gamma+\beta)^2}{\beta}\,. \tag{1.31}$$

Thus, if V exceeds c/γ, then the right hand side of (1.30) is negative and V is forced to decrease.

Tying together the remaining loose ends in Homework 2 and 3, we have altogether proved the following.

Theorem 4. *There exists a bound $B > 0$ depending only on the parameters σ, ρ, and β such that for any solution $Y = (x, y, z)$ to the Lorenz equations* (1.18) *there exists a time $\tau \geq 0$, which may also depend on the initial data, such that*

$$\|Y(t)\| \leq B \tag{1.32}$$

for all $t \geq \tau$. In particular, there is no finite-time blow-up.

Remark 3. This is just one result on the Lorenz equations which will be important later. For more information, see modern textbooks on differential equations, for example [3].

Homework 2. Solve the differential inequality (1.30) explicitly.

Homework 3. Show that $V(t)$ is bounded if and only if $\|Y(t)\|$ is bounded.

Project 2. Repeat the above argument for the van der Pol equation. Note that the first part of the argument, i.e. the proof that the norm of the solution grows at most exponentially, is easy. The hard part is to show that trajectories enter some bounded set independent of their initial value.

2 Nudging

2.1 Direct insertion

We assume that one of the variables, $y(t)$ say, is observable, and that we have complete exact knowledge of its value for all $t \geq 0$. We also assume that the dynamics is accurately represented by the Lorenz model (1.18)—we are not discussing modeling errors here. However, the remaining variables are considered inaccessible; in particular, we assume no knowledge about their initial values. Our task is to use the available information to reconstruct the correct "internal" states $x(t)$ and $z(t)$.

A natural and simple approach is to insert the known time series for y into the equations for x and z, then compute the resulting evolution with arbitrarily chosen initial data. I.e., we solve

$$\dot{\xi} = \sigma\,(y - \xi)\,, \tag{2.1a}$$

$$\dot{\zeta} = \xi y - \beta\zeta\,. \tag{2.1b}$$

with arbitrarily chosen

$$\xi(0) = \xi_0\,, \tag{2.1c}$$

$$\zeta(0) = \zeta_0\,. \tag{2.1d}$$

For convenience, we assume that the initial state is already within the finite bounded region afforded by Theorem 4. This assumption is not essential, but simplifies the discussion in several places.

If we initialized (2.1) with the exact (but unknown) initial data $\xi(0) = x_0$ and $\zeta(0) = z_0$, we would replicate our reference solution exactly. For other initial data we hope that our reconstructions ξ and ζ will at least converge to the reference solution. The numerical experiment shown in Figure 5 suggests that this is indeed so.

We begin our analysis by setting $u = x - \xi$, $v = y - \eta$, and $w = z - \zeta$. (The reconstruction η of the y-component does not appear in system (2.1),

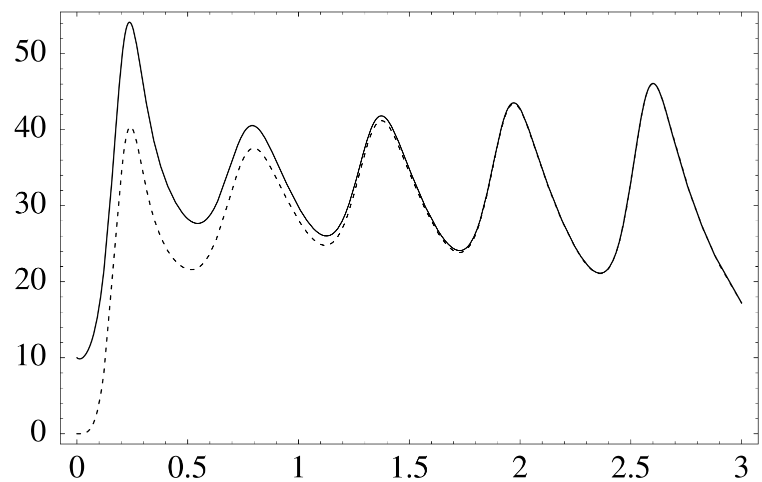

Figure 5: Comparison of the z-component of the Lorenz solution (solid curve) to the solution of the "nudged" equations (2.1) with y taken as the observable (dashed curve). The parameters are as in Figure 3, and the nudged equation is initialized with $\xi_0 = \zeta_0 = 0$.

but is convenient for future use.) By direct computation, using the Lorenz equation (1.18) and the "nudged" system (2.1), we find

$$\dot{u} = \dot{x} - \dot{\xi} = \sigma\,(y - x) - \sigma\,(y - \xi) = -\sigma u\,, \tag{2.2a}$$

$$\dot{w} = \dot{z} - \dot{\zeta} = xy - \beta\,z - (\xi y - \beta \zeta) = yu - \beta w\,. \tag{2.2b}$$

The first equation does not depend on the second and can be solved immediately,

$$u(t) = \mathrm{e}^{-\sigma t}\,u_0\,. \tag{2.3}$$

Thus, $u(t) \to 0$ exponentially as $t \to \infty$, independent of the initial data. We insert this result into (2.2b), and write

$$\dot{w} + \beta w = y\,\mathrm{e}^{-\sigma t}\,u_0\,. \tag{2.4}$$

This equation can be solved after multiplication with a so-called *integrating factor*: We observe that the left hand side is the result of a product rule after multiplying both sides with $\exp(\beta t)$. Hence,

$$\frac{\mathrm{d}}{\mathrm{d}t}(\mathrm{e}^{\beta t}\,w) = \mathrm{e}^{\beta t}\,y\,\mathrm{e}^{-\sigma t}\,u_0\,. \tag{2.5}$$

Integration in time gives

$$\mathrm{e}^{\beta s}\,w(s)\Big|_0^t = u_0 \int_0^t \mathrm{e}^{(\beta-\sigma)s}\,y(s)\,\mathrm{d}s\,. \tag{2.6}$$

The integral on the right can obviously not be solved explicitly since we do not even have an explicit expression for y. However, we only need an estimate that will prove convergence. We thus solve for w, take the absolute value, apply the triangle inequality, and note that Theorem 4 guarantees the existence of a bound $|y(t)| \leq B$ for all time:

$$\begin{aligned} |w(t)| &= \left| \mathrm{e}^{-\beta t}\, w_0 + u_0\, \mathrm{e}^{-\beta t} \int_0^t \mathrm{e}^{(\beta-\sigma)s}\, y(s)\, \mathrm{d}s \right| \\ &\leq \mathrm{e}^{-\beta t}\, |w_0| + |u_0|\, \mathrm{e}^{-\beta t} \int_0^t \mathrm{e}^{(\beta-\sigma)s}\, |y(s)|\, \mathrm{d}s \\ &\leq \mathrm{e}^{-\beta t}\, |w_0| + |u_0|\, B\, \mathrm{e}^{-\beta t} \int_0^t \mathrm{e}^{(\beta-\sigma)s}\, \mathrm{d}s \\ &= \mathrm{e}^{-\beta t}\, |w_0| + |u_0|\, B\, \frac{\mathrm{e}^{-\sigma t} - \mathrm{e}^{-\beta t}}{\beta - \sigma} \\ &\to 0 \text{ as } t \to \infty\,. \end{aligned} \tag{2.7}$$

(Note that in the last integration we have considered the case when $\sigma \neq \beta$. If $\sigma = \beta$, an even simpler argument leads to the same conclusion.) We conclude the following.

Theorem 5. *When applying* nudging by direct insertion *to the Lorenz model with observable y,*

$$|x(t) - \xi(t)| \to 0 \qquad \textit{and} \qquad |z(t) - \zeta(t)| \to 0 \tag{2.8}$$

as $t \to \infty$ for any choice of initial values ξ_0 and ζ_0.

We now proceed to see if it is possible to use z as the only observable. The corresponding nudged equations are

$$\dot{\xi} = \sigma\,(\eta - \xi)\,, \tag{2.9a}$$
$$\dot{\eta} = \rho\xi - \eta - \xi z\,, \tag{2.9b}$$

again with unknown initial values ξ_0 and η_0. By direct computation,

$$\dot{u} = \sigma\,(v - u)\,, \tag{2.10a}$$
$$\dot{v} = \rho u - v - uz\,, \tag{2.10b}$$

so that multiplication of these equations with u and v, respectively, gives

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}u^2 = \sigma uv - \sigma\, u^2\,, \tag{2.11a}$$
$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}v^2 = \rho uv - v^2 - uvz\,. \tag{2.11b}$$

To obtain an estimate on the norm of the difference between the nudged and the reference solution, we add the two equations:

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\big(u^2 + v^2\big) = -\sigma\, u^2 - v^2 + (\sigma + \rho - z)\, uv\,. \tag{2.12}$$

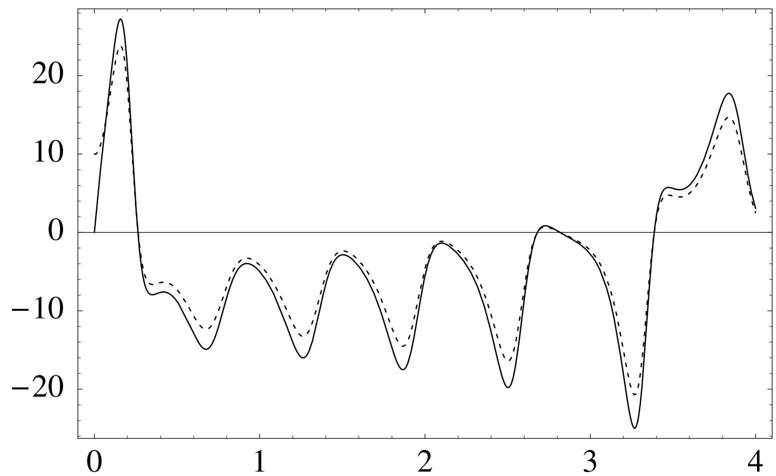

Figure 6: Comparison of the y-component of the Lorenz solution (solid curve) to the solution of the "nudged" equations (2.1) with z taken as the observable (dashed curve). The parameters are as in Figure 3, and the nudged equation is initialized with $\xi_0 = 0$ and $\eta_0 = 10$.

If we want to control the cross term, which is the only term that may allow u and v to increase, the best we can do without additional information is to apply Lemma 3 with $a = (\sigma + \rho - z)\, u$ and $b = v$. We find that

$$(\sigma + \rho - z)\, uv \le \tfrac{1}{2}\,(\sigma + \rho - z)^2\, u^2 + \tfrac{1}{2}\, v^2\,, \tag{2.13}$$

so that

$$\frac{1}{2}\,\frac{\mathrm{d}}{\mathrm{d}t}\big(u^2 + v^2\big) \le \big(\tfrac{1}{2}\,(\sigma + \rho - z)^2 - \sigma\big)\, u^2 - \tfrac{1}{2}\, v^2\,. \tag{2.14}$$

We have no control on the sign of the first term (unless the parameters σ and ρ are very small, and we can show that the corresponding bound on z given by Theorem 4 is also small). Thus, if u and v have the same sign, (2.12) seems to indicate that the difference between nudged and reference trajectory increases rather than decreases, while for u and v of opposite sign, the difference decreases as desired. Note that we encountered the same type of difficulty when we tried to prove boundedness of trajectories in Section 1.4; see equation (1.20). The modification of the argument we proposed there would not suffice here as it yields only bounds of finite magnitude, not convergence to zero. However, it may be possible to prove a weaker statement, see Project 3.

The argument above does not prove, strictly speaking, that we cannot use z as the only observable. It demonstrates, however, that our usual method of proving convergence fails. And indeed, the numerical experiment shown in Figure 6 verifies our suspicion.

Homework 4. Prove that nudging by direct insertion with x as the observable also works.

Project 3. Figure 6 suggests that although the two curves do not converge, they remain within a small distance of each other. Can you prove that this distance is significantly smaller than the typical amplitude of the solution, e.g. as measured by the bound B from Theorem 4?

2.2 Time-discrete data

Often, we have to assimilate data which is only available at discrete times. There are two ways to proceed. First, we could find a function which interpolates the given data points and perform continuous nudging as described previously using the interpolant. In this case, we would need to analyze the interpolation error, a task similar to analyzing measurement errors. We will not proceed further in this direction here. Second, we could solve the full system in between measurements, re-initializing with new data when it becomes available. We will study this procedure in detail, again in the context of the Lorenz equations.

We assume that $x(t)$ is observable at discrete, equidistant times $t_n = nh$, where $n = 0, 1, \ldots$ and h denotes the distance between the sampling points. The functions y and z are assumed to be non-observable. We proceed as follows.

(i) When $t = 0$, set $\xi_0 = x(0)$ and initialize η and ζ arbitrarily.

(ii) Let ξ, η, and ζ evolve from time t to time $t + h$ according to the exact Lorenz equations

$$\dot{\xi} = \sigma\,(\eta - \xi)\,, \tag{2.15a}$$

$$\dot{\eta} = \rho\xi - \eta - \xi\zeta\,, \tag{2.15b}$$

$$\dot{\zeta} = \xi\eta - \beta\zeta\,. \tag{2.15c}$$

(iii) At time $t+h$, keep the values of η and ζ, but set ξ to the new observation $x(t+h)$. Go to step (ii) and repeat with new initial time $t+h$.

Again, we ask the question if, or under which conditions, do the reconstructions ξ, η, and ζ converge to the correct internal states. Numerical experimentation indicates that for small enough step sizes, convergence is achieved (Figure 7), but if the step size grows too large, convergence fails (Figure 8).

We will now analyze a single step of this procedure and prove that for h small enough the difference to the reference solution reduces by at least a constant factor smaller than one. As in the previous section, we compute, using (1.18) and (2.15),

$$\dot{u} = \sigma\,(v - u)\,, \tag{2.16a}$$

$$\dot{v} = \rho u - v - xz + \xi\zeta = \rho u - v - xw - u\zeta\,, \tag{2.16b}$$

$$\dot{w} = xy - \xi\eta - \beta w = xv - u\eta - \beta w\,. \tag{2.16c}$$

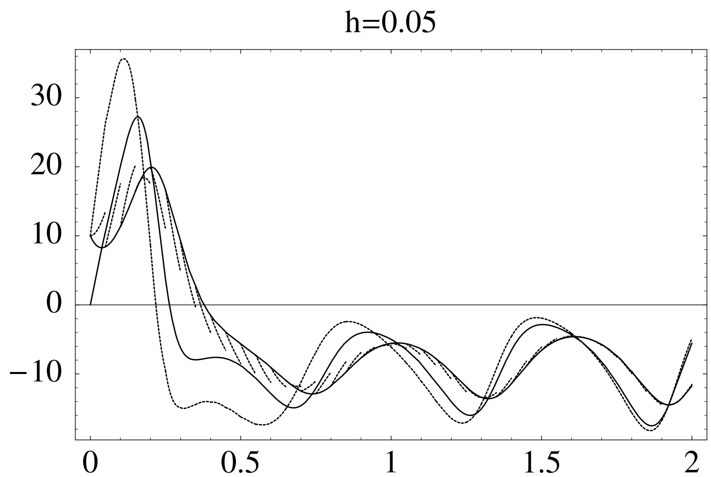

Figure 7: Discrete nudging. Shown are the observable x and internal state y with solid lines, and their reconstructions ξ and η with dashed lines. ξ is clearly distinguished by its discontinuity through re-initialization at the beginning of each step. The parameters are as in Figure 3, and the internal states are initialized with $\eta_0 = 10$ and $\zeta_0 = 0$.

Multiplying these equations with u, v, and w, respectively, we find

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}u^2 = \sigma uv - \sigma\, u^2\,, \tag{2.17a}$$

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}v^2 = \rho uv - v^2 - xvw - uv\zeta\,, \tag{2.17b}$$

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}w^2 = xvw - uw\eta - \beta\, w^2\,. \tag{2.17c}$$

Setting $V^2 = v^2 + w^2$, we obtain by adding (2.17b) and (2.17c) that

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}V^2 = (\rho - \zeta)\, uv - v^2 - uw\eta - \beta\, w^2\,. \tag{2.18}$$

We estimate the cross term in (2.17a) using Lemma 3 with $a = u$ and $b = v$. Similarly, the first cross term in (2.18) is estimated by Lemma 3 with $a = (\rho - \zeta)\, u$ and $b = v$; for the second cross term in (2.18) we take $a = \eta u/\sqrt{\beta}$ and $b = \sqrt{\beta}\, w$. Altogether, we obtain the system of differential inequalities

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}u^2 \le \frac{1}{2}\,\sigma\, v^2 - \frac{1}{2}\,\sigma\, u^2\,, \tag{2.19a}$$

$$\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}V^2 \le \frac{1}{2}\,(\rho - \zeta)^2\, u^2 - \frac{1}{2}\, v^2 + \frac{1}{2}\,\frac{\eta^2}{\beta}\, u^2 - \frac{1}{2}\,\beta\, w^2\,. \tag{2.19b}$$

Theorem 4 provides upper bounds on η and ζ, so that $(\rho - \zeta)^2 \le (\rho + B)^2$ and $\eta^2 \le B^2$. Setting

$$c = \max\{\sigma, (\rho + B)^2 + B^2/\beta\} \tag{2.20}$$

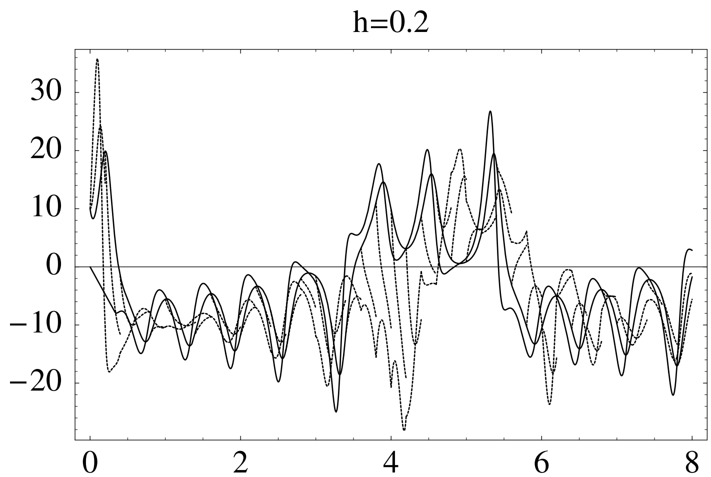

Figure 8: Non-convergence for large step size h. The setting is the same as in Figure 7, but the step size is four times larger.

and

$$\varepsilon = \min\{1, \sigma, \beta\} \tag{2.21}$$

we see that u and V satisfy the system of differential inequalities

$$\frac{\mathrm{d}}{\mathrm{d}t} u^2 \le c\, V^2 - \varepsilon\, u^2 \,, \tag{2.22a}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} V^2 \le c\, u^2 - \varepsilon\, V^2 \,. \tag{2.22b}$$

For simplicity, we assume $c = 1$; the general case can be recovered through an appropriate rescaling of time. We will first solve this system as a differential equation, and then discuss how this solution provides bounds on the behavior of functions satisfying the differential inequalities (2.22). Setting

$$U = \begin{pmatrix} u^2 \\ V^2 \end{pmatrix} \qquad \text{and} \qquad A = \begin{pmatrix} -\varepsilon & 1 \\ 1 & -\varepsilon \end{pmatrix} , \tag{2.23}$$

we now have to solve the linear system

$$\dot{U} = A\, U \,. \tag{2.24}$$

Note that A has the diagonalization $A = SDS^{-1}$ with

$$D = \begin{pmatrix} 1-\varepsilon & 0 \\ 0 & -1-\varepsilon \end{pmatrix} \qquad \text{and} \qquad S = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = S^{-1} \,, \tag{2.25}$$

so that

$$
\begin{aligned}
\mathrm{e}^{At} = S\,\mathrm{e}^{Dt}\,S^{-1} &= S \begin{pmatrix} \mathrm{e}^{(1-\varepsilon)t} & 0 \\ 0 & \mathrm{e}^{-(1+\varepsilon)t} \end{pmatrix} S^{-1} \\
&= \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \mathrm{e}^{(1-\varepsilon)t} & \mathrm{e}^{(1-\varepsilon)t} \\ \mathrm{e}^{-(1+\varepsilon)t} & -\mathrm{e}^{-(1+\varepsilon)t} \end{pmatrix} \\
&= \frac{1}{2} \begin{pmatrix} \mathrm{e}^{(1-\varepsilon)t} + \mathrm{e}^{-(1+\varepsilon)t} & \mathrm{e}^{(1-\varepsilon)t} - \mathrm{e}^{-(1+\varepsilon)t} \\ \mathrm{e}^{(1-\varepsilon)t} - \mathrm{e}^{-(1+\varepsilon)t} & \mathrm{e}^{(1-\varepsilon)t} + \mathrm{e}^{-(1+\varepsilon)t} \end{pmatrix} \qquad (2.26)
\end{aligned}
$$

Since we are initializing ξ exactly, $u(0) = 0$, hence $U(0) = (0, V_0^2)$, so that the solution to (2.24) is given by

$$
U(t) = \mathrm{e}^{At}\,U(0) = \frac{1}{2} \begin{pmatrix} \mathrm{e}^{(1-\varepsilon)t} - \mathrm{e}^{-(1+\varepsilon)t} \\ \mathrm{e}^{(1-\varepsilon)t} + \mathrm{e}^{-(1+\varepsilon)t} \end{pmatrix} V_0^2 \qquad (2.27)
$$

and, in particular,

$$
V^2(h) = \tfrac{1}{2} \left(\mathrm{e}^{(1-\varepsilon)h} + \mathrm{e}^{-(1+\varepsilon)h} \right) V_0^2 \,. \qquad (2.28)
$$

It is easy to see that the prefactor on the right is bounded by a constant $\mu < 1$ for h not too large (see Homework 5), so that, after n iterations of the argument, our error measure has decreased to

$$
V(nh) \leq \mu^{n/2}\,V_0 \,. \qquad (2.29)
$$

This bound converges to zero as $n \to \infty$.

It can be shown that the result of this computation truly provides a bound on u and V so long as V is decreasing (see Project 4), even though we had replaced the inequalities in (2.22) by equalities. Thus, we have altogether proved the following.

Theorem 6. *There exists a maximum step size H so that nudging the Lorenz model with discrete data in the x variable at times $t_n = nh$ with $h \in [0, H]$ yields an approximation (ξ, η, ζ) that converges to (x, y, z) as $n \to \infty$.*

Remark 4. Our proof gives a sufficient condition for convergence. Numerical experiments, on the other hand, indicate that even when h is much larger than the bounds given above, convergence still occurs (Figure 9). The zoom shown in Figure 10 shows that this is not due to "being too generous" in our estimates when analyzing a single step: There are clearly steps in which the computed error does grow. Convergence in this situation is thus a consequence of subtle averaging over many steps, and a rigorous analysis requires completely different techniques. The case when Theorem 6 truly applies is shown in Figure 11: the error decreases monotonically from step to step, although it may increase within a single step.

Homework 5. Prove that such $\mu < 1$ exists as claimed in the text, and give an estimate in terms of ε of the maximum possible value for h.

Project 4. Give an argument why it was possible to take the solution of (2.24) as an upper bound for u and V satisfying the differential *inequalities* (2.22).

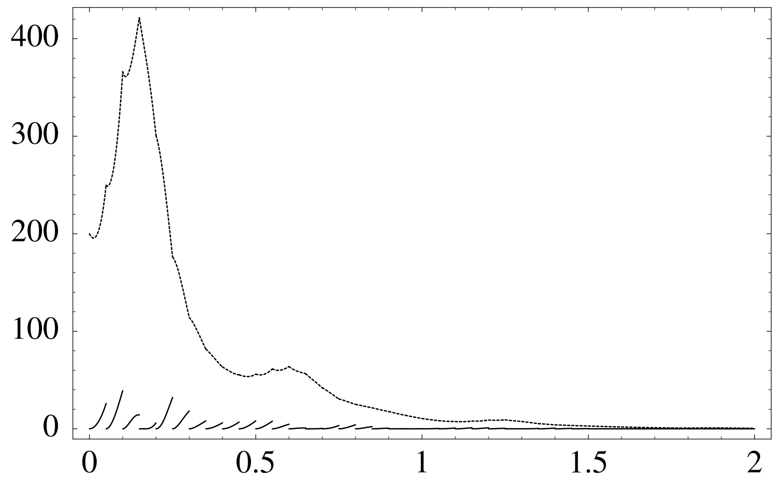

Figure 9: Error measures u^2 (solid line) and V^2 (dashed line) for the solutions plotted in Figure 7.

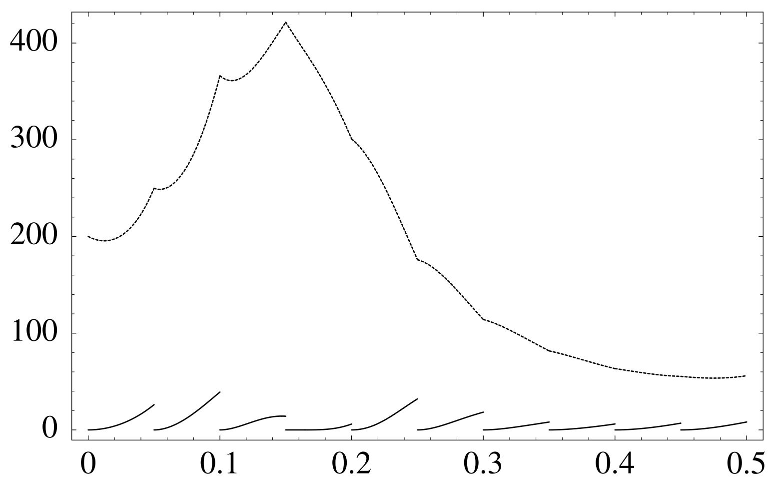

Figure 10: Zoom into the first few steps shown in Figure 9. One can clearly see that in each nudging interval the error initially decreases, but the step size h is too large to guarantee a decrease in error up to the end of the interval.

2.3 Newtonian relaxation

In practical applications, nudging is often not done by direct insertion, but by introducing additional forces into the full system that try to "relax" the system

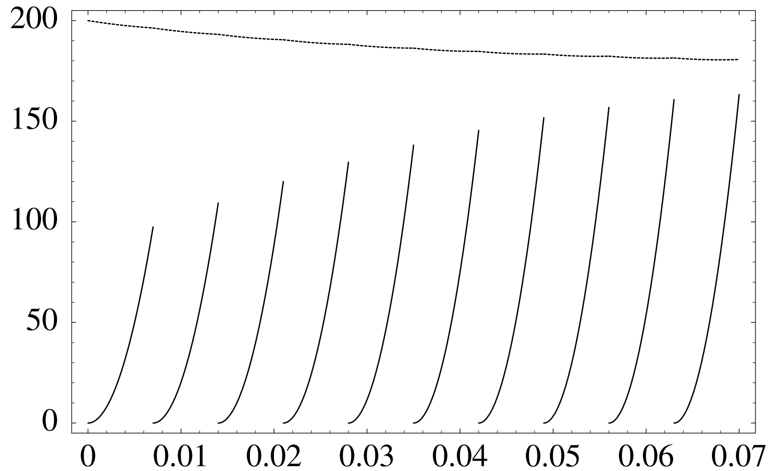

Figure 11: Discrete nudging with tiny step size. In this case, Theorem 6 applies and the error measure V^2 (dashed line) decreases across each nudging interval. The error measure u^2 (solid line) is so small that it is magnified by a factor of 200.

to the desired trajectory. For example, when nudging in x, we may solve

$$\dot{\xi} = \sigma\,(\eta - \xi) + \frac{x - \xi}{\tau}\,, \tag{2.30a}$$

$$\dot{\eta} = \rho\xi - \eta - \xi\zeta\,, \tag{2.30b}$$

$$\dot{\zeta} = \xi\eta - \beta\zeta\,. \tag{2.30c}$$

Project 5. Analyze this procedure following the proofs in the previous two sections. By varying τ, one can trade off speed of convergence versus robustness with regard to noise in the data. Explore this claim by computer experiment and/or analysis.

3 Comments

Nudging is arguably the simplest form of sequential data assimilation. It is easy to implement, fast, and, depending on the particular situation, relatively robust and transparent. However, it has also some distinct disadvantages. First, nudging needs relatively large amounts of data. It cannot distribute the reconstruction error uniformly across the time interval of interest, and is therefore suboptimal in the usage of the supplied data. Second, it relies fundamentally on dissipativity, i.e. the nonconservation of energy, in the equation. However, one sometimes would like to assimilate data for systems that are not dissipative, or are only very slightly so.

An alternative to nudging is *variational data assimilation*, in which some measure for the error, for example the function U of Section 2.2, is actively minimized. Such a procedure is naturally more expensive and more difficult to implement. On the other hand, variational data assimilation performs much better when data is sparse and the equations are not very dissipative. It also involves interesting mathematics, but this story will be told another time. A few entry points for further reading, available online, are [1, 2, 6].

Project 6. Expand this critique of nudging, and survey methods of variational data assimilation.

References

[1] M.S. Roulston, *Data assimilation in chaotic systems*, `http://gfd.whoi.edu/proceedings/1998/PDF/rep%5Froulston.pdf`

[2] T.W. Schlatter, *Four-dimensional data assimilation*, in "The Analysis of Observations with Applications in Atmospheric Science", `http://www.asp.ucar.edu/colloquium/1992/notes/`

[3] F. Verhulst, "Nonlinear Differential Equations and Dynamical Systems", Springer-Verlag, 1996.

[4] Eric W. Weisstein, *Lorenz Attractor*, MathWorld, Wolfram Research, `http://mathworld.wolfram.com/LorenzAttractor.html`

[5] *Lorenz equations*, The Math Forum, Drexel University, `http://mathforum.org/library/drmath/view/61791.html`

[6] *Understanding data assimilation: how models create their initial conditions*, University Corporation for Atmospheric Research, `http://meted.ucar.edu/nwp/pcu1/ic6/`