



WOODS HOLE OCEANOGRAPHIC INSTITUTION
DEPARTMENT OF PHYSICAL OCEANOGRAPHY
MS #29, WOODS HOLE, MASSACHUSETTS, 02543
TEL: (508) 289-2801, ggebbie@whoi.edu
Geoffrey Gebbie

January 6, 2017

Nonlinear Processes in Geophysics Discussions
European Geophysical Union

Dear Dr. Zoltan Toth, Editor of *Nonlinear Processes in Geophysics Discussions*,

Please find under this cover a revision of the manuscript entitled, "Controllability, not chaos, key criterion for ocean state estimation," by Geoffrey Gebbie and Tsung-Lin Hsieh. We have taken into account the constructive and thorough suggestions from the three reviewers, and I hope you agree that it has led to an improved and clarified work.

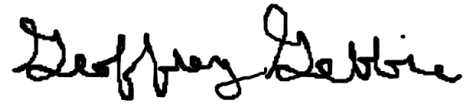
We have paid particular attention to Reviewer #1's comments regarding the context of our work. We have added new paragraphs in the discussion of controllability and observability that better ties our work to the decades of previous research. New paragraphs have also been included in the Introduction to better describe our motivation and aims. Much time and effort has been spent to develop the Lagrange multiplier method in real-world scenarios, yet it is unclear whether this method should be applied to eddy-resolving models and how long the time window should be. Rather than developing a novel data assimilation technique, we implement diagnostics for stability and controllability that help us understand when the method will be successful. Unlike previous toy models, we suggest a greater range of applicability for the Lagrange multiplier method than is commonly thought, which will interest the data assimilation community. In fact, our results with a chaotic pendulum toy model suggest that there is no fundamental obstacle to ocean state estimation with eddy-resolving, highly-nonlinear models.

We have also paid particular attention to two reviewers comments about the chosen experimental system. To make the setup more realistic, we have performed additional experiments where the first-guess of external forcing is not perfect. This information is in a new Section (3.5) and includes a new Figure (8). These additional findings support our original interpretations. Certainly, research with a more complex geophysical model would be fruitful. Here, we emphasize that for the Lagrange multiplier method to be successful in state-of-the-art ocean models, two major issues need to be addressed: (1) the high dimensionality of the forward model and estimation problem, and (2) the nonlinearity of ocean models at increasingly fine resolution. Issue (1) has been overcome by groups such as the ECCO Consortium. Here we focus on (2) and isolate the effect of nonlinearity by choosing this particular toy model.

Thank you again for considering this work for publication. The reviewers have provided a fresh perspective that, we believe, has led to an improved manuscript. A detailed point-by-point response to the reviewers

follows.

Sincerely,

A handwritten signature in black ink that reads "Geoffrey Gebbie". The signature is written in a cursive, slightly slanted style.

Geoffrey (Jake) Gebbie,
on behalf of Tsung-Lin Hsieh

Reviewer #1

- *This manuscript contains a description of the results of a series of experiments in which 4DVAR was used to assimilate simulated data from a nonlinear system corresponding to the damped driven pendulum in a chaotic parameter regime. This example differs from most other examples in the literature of data assimilation in strongly nonlinear systems in that it is a non-autonomous system, unlike, say, Lorenz (63). With a few reservations, the example is fairly well worked out. The application of the χ^2 test is particularly noteworthy. The basic results are worth publishing in some form.*

We thank the reviewer for noting the care we took in this analysis, especially the χ^2 posterior test, and for noting that the key results are worth publishing. We address some of your reservations in this point-by-point response.

- *The authors never state their model system explicitly. It is not (1). The system with which they are actually working differs from (1) in that it has a white noise term with variance S_f (see (5)) added to the right hand side. The distinction is not trivial. I assume that the reference solution in their twin experiments is the stochastic system with the stochastic term set to zero. The effect of adding the unknown stochastic term is to increase the number of degrees of freedom in the control space from two, i.e., the initial conditions in the purely deterministic problem, to the number of time steps taken by the numerical method, which is potentially infinite.*

Thanks to the reviewer for spotting the error in defining the external forcing term. The equation is deterministic and is now stated explicitly in equation (1). The weight matrix S_f in equation (5) is selected in the state estimation process to limit the difference of the improved guess from the first guess, similar to the formulation in Bennett (2002). The revised text reads as follows.

The motion of the forced pendulum is described by the deterministic equation (Baker and Gollub, 1990),

$$\frac{d^2\theta}{dt^2} + \frac{1}{q} \frac{d\theta}{dt} + \frac{g}{l} \sin \theta = f(t), \quad (1)$$

where θ is the displacement angle from vertical, q is a damping coefficient, g is gravitational acceleration, l is the pendulum length, and $f(t)$ is an external forcing term. In turn, the external forcing has a first guess and a perturbation, $f(t) = f_0 + \delta f(t)$, where the first-guess is set to periodic forcing, $f_0(t) = b \cos(\omega_d t)$.

- *The general level of discussion in this manuscript might have been marginally acceptable twenty years ago, when implications of applications of techniques from the engineering world were still being explored, but most of the manuscript is far below the current state of the art.*

It is clear that the motivation and goals of the manuscript need to be made more explicit. The Lagrange multiplier method is popular in oceanography due to automatic adjoint model compilers and strategies to reduce computer memory consumption. Much time and effort has been spent to develop this technique in real-world scenarios, yet it is unclear whether this method should be applied to eddy-resolving models and how long the time window should be. For the Lagrange multiplier method to be successful in state-of-the-art ocean models, two major issues need to be addressed: (1) the high dimensionality of the forward model and estimation problem, and (2) the nonlinearity of ocean models at increasingly fine resolution. Issue (1) has been overcome by groups such as the ECCO Consortium. Here we focus on (2). It is true that issues may arise by the combined effect of (1) and (2), but first we attempt to isolate the effect of nonlinearity.

With this problem in mind, it is logical to find a numerical model that can be thoroughly understood and one that is highly nonlinear. It is not the goal of this manuscript to use a state-of-the-art numerical model. We believe that these expectations should be set at the outset, so we include the following in the Introduction.

Because the effect of nonlinearity is seen as the major roadblock for application of the Lagrange multiplier method, we isolate this effect by choosing a model that is highly nonlinear but low-dimensional: the forced, chaotic pendulum (Section 2). Toy models are worth revisiting because the dynamics are comparatively simple to understand, and they have strongly influenced when the Lagrange multiplier method has been deployed to realistic ocean problems. We will show that previous toy models have sometimes been misinterpreted.

We now also emphasize upfront that the development of a new state-of-the-art data assimilation technique is not the goal of this work, either. Instead, we wish to evaluate the current use of the Lagrange multiplier method. Now, the Introduction makes this explicit.

Rather than developing a new state-of-the-art data assimilation technique, we proceed by taking the existing Lagrange multiplier method and developing diagnostics regarding when and why it succeeds or fails, as evaluated by the ability to fit observations. Relative to the initialization problem, the prospects for a successful state estimate are shown to be improved in the boundary control problem, even if one uses a highly nonlinear model such as the forced, chaotic pendulum (Section 3).

The typical criterion for successful state estimation has been the stability to initial perturbation. In the manuscript, we provide a counterexample showing that state estimation can be successful for an unstable system when it is controllable. This is one novel result we are reporting, and the previous works suggested by the reviewers have not already made this point, nor do they appear to contradict it.

- *Studies of chaotic systems forced by white noise have appeared in a number of places in the literature. One example can be found in a paper by Tziperman from the early 90s.*

We now mention Tziperman's work on chaotic systems as a motivating factor in using our toy model.

Toy models are worth revisiting because the dynamics are comparatively simple to understand, the nonlinear coupling to periodic forcing has been shown to be important in atmosphere-ocean dynamics (e.g., *Tziperman et al.*, 1994), and these models have strongly influenced when the Lagrange multiplier method has been deployed to realistic ocean problems.

- *There is nothing novel about writing the 4DVAR cost function in terms of a Lagrange multiplier. The use of Lagrange multipliers in variational formulations of estimation and control problems has been in the engineering textbooks since the 70s, and appeared in the early work of Thacker in the ocean modeling literature. In the present context, in which the task is to estimate an unknown stochastic forcing function, the Lagrange multiplier formulation is valid, but the same Euler-Lagrange equations result from equivalent cost function formulations without Lagrange multipliers, see, e.g., the text by Kalnay or either of the books by Bennett, as well as many of the reviews in the literature.*

The reviewer's point that the Lagrange multiplier method is just another way to express a minimization problem is in line with the point of our Section 4.1. In that section, we suggest that the criteria for successful use of the Kalman Filter/Smoothing, which also minimizes the same cost function, is the same as that for the Lagrange multiplier method. In particular, Sec. 4.1 states the following.

Our results suggest that the equivalence of the Kalman filter/smoothing and Lagrange multiplier method may be extended to nonlinear problems, thus explaining why the chaotic estimation problem may be solved by the Lagrange multiplier method.

- *The authors should note that the estimation problem is the dual of the control problem. General questions of linear controllability and observability are dealt with in engineering textbooks. This topic has been well worked out in the context of models of the ocean and atmosphere in the work of S. E. Cohn in the late 80s and early 90s. The question of nonlinear observability is very complex. There was a book by Casti on the subject published some time ago.*

We thank the reviewer for the excellent suggestions for further references. In addition, we now point out the duality of the estimation and the control problem. As pointed out by the reviewer, observability and controllability conditions for nonlinear state estimation are difficult problems. Results for certain nonlinear systems are found in the book by *Casti* (1985) (which evolves from the review paper of *Casti* (1982)). For general nonlinear estimation problems, the down-gradient-based iterative optimization is likely one of the best methods, and we have shown the relevance of controllability to the iterative process in our example. We include the following new paragraph in the Discussion.

To recover the true trajectory of a system, observability is also important, as the estimation problem is the dual of the control problem (*Fukumori et al.*, 1993; *Marchal*, 2014). For the linear problem, *Cohn and Dee* (1988) showed that complete observability implies asymptotic stability of the Kalman filter/smoothen. Defining observability and controllability conditions for nonlinear state estimation problems is difficult *Casti* (1985). In practice, the important criterion is ability to solve equation (10). Strictly speaking, the solution criteria will therefore depend upon both the controllability matrix, \mathbf{C} , and the observational matrix, \mathbf{E} , which combines the issues of observability and controllability. Here, we suggest the operational definition that a system is effectively controllable when the solution to (10), generalized to multiple observations, exists.

- *The question of dealing with underdetermination has been discussed extensively in the literature. Solutions to underdetermined problems are not, in general, unique. The problem, in practice, is the fact that minimizing the cost function (5) involves searching a space of corrections that is potentially infinite. The highly irregular reconstructed forcing shown in the bottom panel of figure 4 is most likely one of an enormous number of minimizers of (5). There are almost certainly many others that will minimize the cost function, some smoother, many even more irregular.*

The reviewer's point about the underdetermined nature of the problem is consistent with our discussion in Sec. 3.3, where we acknowledge that the solution is not unique, but we focus on finding any acceptable fit. We view the problem as having two clear steps. It is a first step to find any solution that acceptably fits the data. Only then can we proceed to investigate the uniqueness of the solution. In real-world situations, the first step may be the only one that is practical.

- *Bennett showed that, in the linear problem, one solution can be found by choosing a correction to the first guess that lies in an N_y dimensional space spanned by representer functions, where N_y is the number of observations. This solution corresponds to the Moore-Penrose inverse. Arguments as to why that solution should be preferred over others are the stuff of textbooks.*

It is worth clarifying that the external forcing had been treated as a controllable parameter in many works in the literature. In the Introduction, we now state the following.

As has been documented in detail by many authors including the textbook of *Bennett* (1992), the ocean state estimation problem is better described as a time-variable boundary value problem because synoptic atmospheric variability acts as an external forcing on the ocean (Section 2). Given our relatively uncertain knowledge regarding air-sea fluxes, the ocean state estimation is rightfully considered a time-variable boundary value problem where both the initial conditions and boundary conditions must be found. For example, *Bennett* (2002) described an estimation method for the external forcing, initial and boundary conditions that solves the Euler-Lagrange equations for a linear model.

- *Similar practical results can be had without explicit calculation of representers. In practical problems in modeling the ocean and atmosphere, the correction to the forcing function lies in a space of enormous dimension, so it is common to precondition the search for a cost function minimizer. This effectively reduces the dimension of the control space by choosing corrections to be a linear combination of singular vectors of the error covariance matrix. This approach is documented in the work*

of A. Lorenz and O. Talagrand. In the present problem, it might be reasonable to impose nontrivial temporal correlation on the forcing correction, which might have the effect of limiting the spectrum of the correction and thus ruling out irregular forcing corrections like that shown in figure 4.

As the reviewer suggests, temporal correlations in the forcing field can be imposed through the use of nondiagonal weighting matrix, \mathbf{S}_f , in the cost function. The revised manuscript now describes how temporal correlations have been enforced in our analysis. The following material has been added to Sec. 3.4.

We investigate the effect of a decrease in the number of controls by redefining the external forcing control perturbation. For N_u forcing controls, we define,

$$f(t) = f_0(t) + \Gamma(t) \begin{pmatrix} \delta f(0) \\ \delta f(T/N_u) \\ \delta f(2T/N_u) \\ \vdots \\ \delta f(T) \end{pmatrix}, \quad (2)$$

where $\Gamma(t)$ is a matrix that performs linear interpolation in time, and $\delta f(t)$ is only defined at N_u control times. This formulation enforces some temporal correlation in the external forcing. Alternatively, this could be accomplished using a nondiagonal weighting matrix, \mathbf{S}_f .

- *The authors have a choice. They can simply report on the results of their twin experiments on their nonautonomous system and eliminate nearly all of the discussion, or they can go back over twenty or twenty five years of literature and rewrite the discussion to make it a meaningful contribution to the current state of the art.*

We have taken seriously the reviewer's suggestions to place our work in the greater context of the published literature. Major revisions include new paragraphs in the Introduction regarding the motivation and aims, more information about the background of controllability that places our work in a broader context, and a new figure using an imperfect model that recreates a more realistic scenario.

Reviewer #2

- *Main points: The perfect model and observations may not be sufficient in supporting the conclusions reached in this manuscript. Even though the reviewer agrees with the authors 4DVAR has its potential in oceanic state estimation, their case is simply too perfect for convincing the readers, particularly those from the non-variational analysis community, like EnKF or 4DVar. The reviewer suggests more realistic experiments and recommend this manuscript for major revision.*

Thanks to the reviewer for the interesting and helpful perspective. The experimental setup may not be as idealized as originally thought. For example, the synthetic observations generated in this study are not perfect, and we detail their generation in the response to the next point below.

Regarding the perfect model assumption, our equations have been formulated in analogy to the ECCO (Estimating the Climate and Circulation of the Ocean) state estimation equations (Stammer *et al.*, 2002), as our goal is to develop diagnostics regarding when and why the Lagrange multiplier method succeeds or fails. In the ECCO formulation, the ocean model equations are typically treated as perfect in the ocean interior, and errors are permitted in the surface gridcells with air-sea forcing. Here, we permit these errors and attribute them to errors in the external forcing.

One aspect of the analysis that we have improved is the first-guess of the forcing field. In a case study where the first-guess of the forcing is zero, the results are similar to the original case. This is reported in a new Section 3.5 and a new Figure 8. The new section follows.

The previous examples in Section 3 proceed with prior information that the forcing is periodic with an accurate magnitude and phase. A good analogy is the regular forcing of solar insolation on the ocean surface. Here, we test the performance of the Lagrange multiplier method with inaccurate prior information about the forcing, as is a more realistic analogy to the uncertainty of air-sea fluxes. In particular, our first guess of the forcing, $f_0(t)$, is systematically biased by decreasing b from 1.5 to 0.75 rad s⁻². The trajectory driven by inaccurate forcing is no worse than the previous cases with accurate forcing due to the dominance of the chaotic dynamics of system (Figure 8). Using the same observations as shown in Figure 3, we find that the chaotic pendulum trajectory is tracked over multiple nonlinear timescales despite this more stringent test. In this case, however, the forcing estimate still contains errors relative to the true forcing calculated with $b = 1.5$ rad s², and some high-frequency structures remain in $f(t)$ (see “improved first guess” in bottom panel, Figure 8). If instead the Lagrange multiplier method is started from the standard first guess, a smoother and more accurate estimate of the forcing is obtained at the expense of not fitting the data as well (see “final estimate” in bottom panel). Any remaining irregular structures can be handled by imposing temporal correlations as was done in Section 3.4. If such measures are not taken, the investigator must take care to decide what elements of the forcing represent true variability and which are compensating for model error. In our simple system of equations, model errors and forcing errors are mathematically equivalent. In state estimates with eddy-resolving GCMs, however, smallscale forcing variability is found near oceanic fronts and the investigator must determine on a case-by-case basis to what extent it reflects real variability.

Convincing the non-variational analysis community to adopt the Lagrange multiplier method would be a challenging task, but outside the goals of our work. It is clear that the motivation and goals of the manuscript need to be made more explicit. The Lagrange multiplier method is popular in oceanography due to automatic adjoint model compilers and strategies to reduce computer memory consumption. Much time and effort has been spent to develop this technique in real-world scenarios, yet it is unclear whether this method should be applied to eddy-resolving models and how long the time window should be. For the Lagrange multiplier method to be successful in state-of-the-art ocean models, two major issues need to be addressed: (1) the high dimensionality of the forward model and estimation problem, and (2) the nonlinearity of ocean models at increasingly fine resolution. Issue

(1) has been overcome by groups such as the ECCO Consortium. Here we focus on (2). It is true that issues may arise by the combined effect of (1) and (2), but first we attempt to isolate the effect of nonlinearity.

With this problem in mind, it is logical to find a numerical model that can be thoroughly understood and one that is highly nonlinear. It is not the goal of this manuscript to use a state-of-the-art numerical model. We believe that these expectations should be set at the outset, so we include the following in the Introduction.

Because the effect of nonlinearity is seen as the major roadblock for application of the Lagrange multiplier method, we isolate this effect by choosing a model that is highly nonlinear but low-dimensional: the forced, chaotic pendulum (Section 2). Toy models are worth revisiting because the dynamics are comparatively simple to understand, and they have strongly influenced when the Lagrange multiplier method has been deployed to realistic ocean problems. We will show that previous toy models have sometimes been misinterpreted.

We now also emphasize upfront that the development of a new state-of-the-art data assimilation technique is not the goal of this work, either. Instead, we wish to evaluate the current use of the Lagrange multiplier method. Now, the Introduction makes this explicit.

Rather than developing a new state-of-the-art data assimilation technique, we proceed by taking the existing Lagrange multiplier method and developing diagnostics regarding when and why it succeeds or fails, as evaluated by the ability to fit observations. Relative to the initialization problem, the prospects for a successful state estimate are shown to be improved in the boundary control problem, even if one uses a highly nonlinear model such as the forced, chaotic pendulum (Section 3).

Thus, a practical goal of this work is to convince those groups that already use the Lagrange multiplier method to reconsider the range of scenarios in which they apply the method.

- 1. *What are the values of σ_θ ? From the true solution? I am afraid if ideal observations are used, it does not imply the conclusions made for this ideal model to be useful.*

Stochastic noise is used to generate synthetic data, mimicking the imperfection of ocean observations. This is a common approach that has appeared in ocean state estimation studies such as Tziperman et al. (1992), who used the same iterative adjoint method on a simplified ocean GCM (the momentum equations are balanced and the nonlinear advection is neglected). The manuscript states the following.

We consider an “identical twin” experiment where the true solution is known (solid line, Figure 1), and we observe the pendulum angle episodically through time with normally-distributed random errors of standard deviation, $\sigma_\theta = 0.5$ rad. In most oceanographically-relevant cases, observations have already been collected over some fixed time interval ($0 \leq t \leq T$). Here, observations, $y(t)$, are taken at a set of N_y evenly-spaced times with an time interval of $\Delta t_y = T/(N_y - 1)$.

- 2. *Page 4, line 23, The statement of . . . the quantity inside curl brackets vanishes is not generally true. To do so, there needs an additional term, penalizing the constraint.*

The terms in the curly bracket vanish by definition of our time-stepping model in equation (2). There was an error in defining the external forcing term which may have caused confusion. The equation is deterministic and is now stated explicitly in equation (1). The weight matrix S_f in equation (5) is selected in the state estimation process to limit the difference of the improved guess from the first guess, similar to the formulation in Bennett (2002). The revised text reads as follows.

The motion of the forced pendulum is described by the deterministic equation (*Baker and*

Gollub, 1990),

$$\frac{d^2\theta}{dt^2} + \frac{1}{q} \frac{d\theta}{dt} + \frac{g}{l} \sin \theta = f(t), \quad (3)$$

where θ is the displacement angle from vertical, q is a damping coefficient, g is gravitational acceleration, l is the pendulum length, and $f(t)$ is an external forcing term. In turn, the external forcing has a first guess and a perturbation, $f(t) = f_0 + \delta f(t)$, where the first-guess is set to periodic forcing, $f_0(t) = b \cos(\omega_d t)$.

- 3. Page 5, line 23. *That is where the problem is that such as observation and prior knowledge and freely-running forward model are not enough.*

We agree with the reviewer that the first-guess may not always be sufficient to track a chaotic system. For this reason, we implement a $\chi - 2$ test, detailed in the next point-by-point response, that diagnoses the likelihood of success or failure.

- 4. *For solving a global minimization problem of (6), the first guess is crucial as the authors stated in page 4 line 25-26. However, the improved initial guess of their work presented in Section 2.4 cannot guarantee the initial guess is good enough for converging to the global minimizer. The authors should at least present convincing arguments of why they believe their improved initial guess could reach their goal. To the reviewer, the improved initial guess may fall into the same valley as the original initial guess.*

The reviewer correctly states that there is no guarantee that a solution will be a global minimizer. We discuss the underdetermined nature of the problem in Sec. 3.3, where we acknowledge that the solution is not unique. Instead, we focus on finding any acceptable fit. We view the problem as having two clear steps. It is a first step to find any solution that acceptably fits the data. Only then can we proceed to investigate the uniqueness of the solution. In real-world situations, the first step may be the only one that can actually be evaluated.

To consider whether a solution is an acceptable fit, we include Figure 6 which details the size of the cost function for various numbers of controls and observations. By implementing a χ^2 posterior statistical test, we determine the ratio of success to failure for various parameter ranges. After running many trials, we do not guarantee the results for any particular number of observations and controls, but a clear pattern emerges. We suggest that the pattern of Figure 6 is explained by the basic metrics of controllability and observability. rather than the stability of the system. This is one novel result we are reporting, and the previous works suggested by the reviewers have not already made this point, nor do they appear to contradict it.

Reviewer #3

- *This paper investigates the key criterion for ocean state estimate, which is commonly called data assimilation (DA) in the oceanography and meteorological communities. There have been a lot of theoretical research and development over the past 30 years, the number of literature is just too many to list. The Lagrange multiplier method is just another way to express the minimization problem presented in the traditional 4D-Var. This can be found out Andrew Bennetts 1992 book: Inverse Methods in Physical Oceanography. Cambridge Monographs on Mechanics and Applied Mathematics. Cambridge University Press.*

The reviewer is clearly correct in their assessment of the relationship between 4D-VAR and the Lagrange multiplier method. We prefer the notation, “Lagrange multiplier method,” because it will hopefully be understood by scientists in the greater mathematics and physics communities. We include a sentence in the Introduction that shows that these terms are interchangeable.

The Lagrange multiplier method (e.g., *Thacker and Long*, 1988; *Wunsch*, 2010), sometimes called the adjoint method (e.g., *Hall et al.*, 1982; *Tziperman and Thacker*, 1989), “4D-VAR” (e.g., *Courtier et al.*, 1994; *Ferron and Marotzke*, 2003), or variational data assimilation (e.g., *LeDimet and Talagrand*, 1986; *Bonekamp et al.*, 2001; *Bennett*, 2002), is a method that satisfies both of these criteria, unlike the Kalman filter (e.g., *Fukumori and Malanotte-Rizzoli*, 1995) or nudging techniques (e.g., *Malanotte-Rizzoli and Tziperman*, 1995).

The reviewer’s point that the Lagrange multiplier method is just another way to express a minimization problem is in line with the point of our Section 4.1. In that section, we suggest that the criteria for successful use of the Kalman Filter/Smoothing, which also minimizes the same cost function, is the same as that for the Lagrange multiplier method.

- *The reviewer does not agree the statement: the dimensionality of many million state variables is not a fundamental problem. I think both high dimensionality and nonlinearity of ocean models are challenging issues for the ocean prediction and data assimilation. The controllability in this article is very vaguely defined. In fact the importance of boundary condition in ocean state estimation has long been recognized.*

Some of the sentences in the Introduction may have given the impression that we are the first to consider the ocean state estimation problem as a time-variable boundary control problem. Given our long list of references on the topic, this is obviously not true! We hope to strike a better balance by revising the final paragraph of the Introduction as follows.

As has been documented in detail by many authors including the textbook of *Bennett* (1992), the ocean state estimation problem is better described as a time-variable boundary value problem because synoptic atmospheric variability acts as an external forcing on the ocean (Section 2).

The Estimating the Climate and Circulation of the Ocean (ECCO) Consortium has specifically focused on large-scale ocean state estimation with ocean models of coarse-enough resolution that they are essentially linear (e.g., *Stammer*, 1997; *Stammer et al.*, 2002). In this context, *Wunsch and Heimbach* (2007) stated, “The main issue for the oceanographic problem is one of dimension,” in accordance with the reviewer’s comment. The success of the ECCO Consortium, insofar as a solution can be found that fits the ocean data, indicates that the dimensionality of the problem can be overcome and therefore is not a fundamental obstacle. Here we revise the sentence.

Research conducted by the ECCO (Estimating the Circulation and Climate of the Ocean) Consortium (*Stammer et al.*, 2002, 2004) has demonstrated that (1), the dimensionality of many million state variables presents a challenge, but, insofar as a solution can be found that fits the ocean data, it can be overcome and it does not pose a fundamental obstacle.

Observability and controllability conditions for nonlinear state estimation are difficult problems. Re-

sults for certain nonlinear systems are found in the book by *Casti* (1985) (which evolves from the review paper of *Casti* (1982)). For general nonlinear estimation problems, the down-gradient-based iterative optimization is likely one of the best methods, and we have shown the relevance of controllability to the iterative process in our example. We include the following new paragraph in the Discussion that also expands upon observability and controllability.

To recover the true trajectory of a system, observability is also important, as the estimation problem is the dual of the control problem (*Fukumori et al.*, 1993; *Marchal*, 2014). For the linear problem, *Cohn and Dee* (1988) showed that complete observability implies asymptotic stability of the Kalman filter/smoothen. Defining observability and controllability conditions for nonlinear state estimation problems is difficult *Casti* (1985). In practice, the important criterion is ability to solve equation (10). Strictly speaking, the solution criteria will therefore depend upon both the controllability matrix, \mathbf{C} , and the observational matrix, \mathbf{E} , which combines the issues of observability and controllability. Here, we suggest the operational definition that a system is effectively controllable when the solution to (10), generalized to multiple observations, exists.

- *My view is that the system the authors have employed for the study and the observations are too simplified to draw valuable conclusions for the DA research and development communities. The reviewer recommends a major revision to include more realistic models.*

Our experimental setup may not be as simple as the reviewer understood. We note that stochastic noise is used to generate the synthetic observations, mimicking the imperfect nature of ocean observations. This is a common approach that has appeared in ocean state estimation studies such as *Tziperman et al.* (1992).

Regarding the simplified model of this study, it has been explicitly chosen with the motivation and goals of the manuscript in mind. Much time and effort has been spent to develop the Lagrange multiplier method in real-world scenarios, yet it is unclear whether this method should be applied to eddy-resolving models and how long the time window should be. For the Lagrange multiplier method to be successful in state-of-the-art ocean models, two major issues need to be addressed: (1) the high dimensionality of the forward model and estimation problem, and (2) the nonlinearity of ocean models at increasingly fine resolution. Issue (1) has been overcome by groups such as the ECCO Consortium, leading us to focus on (2).

With this problem in mind, it is logical to find a numerical model that can be thoroughly understood and one that is highly nonlinear. It is not the goal of this manuscript to use a state-of-the-art numerical model. We believe that these expectations should be set at the outset, so we include the following in the Introduction.

Because the effect of nonlinearity is seen as the major roadblock for application of the Lagrange multiplier method, we isolate this effect by choosing a model that is highly nonlinear but low-dimensional: the forced, chaotic pendulum (Section 2). Toy models are worth revisiting because the dynamics are comparatively simple to understand, and they have strongly influenced when the Lagrange multiplier method has been deployed to realistic ocean problems. We will show that previous toy models have sometimes been misinterpreted.

We now also emphasize upfront that the development of a new state-of-the-art data assimilation technique is not the goal of this work, either. Instead, we wish to evaluate the current use of the Lagrange multiplier method. Now, the Introduction makes this explicit.

Rather than developing a new state-of-the-art data assimilation technique, we proceed by taking the existing Lagrange multiplier method and developing diagnostics regarding when

and why it succeeds or fails, as evaluated by the ability to fit observations. Relative to the initialization problem, the prospects for a successful state estimate are shown to be improved in the boundary control problem, even if one uses a highly nonlinear model such as the forced, chaotic pendulum (Section 3).

The typical criterion for successful state estimation has been the stability to initial perturbation. In the manuscript, we provide a counterexample showing that state estimation can be successful for an unstable system when it is controllable. This is one novel result we are reporting, and the previous works suggested by the reviewers have not already made this point, nor do they appear to contradict it.

One aspect of the analysis that we have improved is the first-guess of the forcing field. In a case study where the first-guess of the forcing is zero, the results are similar to the original case. This is reported in a new Section 3.5 and a new Figure 8.

Without doubt, it is a worthy goal to demonstrate the importance of controllability for the Lagrange multiplier method with a more complex geophysical model. Based on the results of this manuscript, a thorough test with that type of model is a logical next step, as suggested in Sec. 5.

- *The presentation also needs improvement, there are some sentences that are either not correct or clear to the reader. e.g. Page 2, line 34-35 is not correct.*

We believe lines 34-35 to be grammatically correct as originally formulated.

- *Page 3, equation (1). It is stated that $\omega_d = 2/3$, so the right hand side forcing $f(t)$ is just a cosine function of time, I cannot see two independent variables (ω , θ) in this equation.*

Certainly the angular velocity and displacement are related by a time derivative and are not independent. Those two variables describe the state of the model, and their time tendencies depend on one another, constituting an algebraic system of equations (discrete version of a system of differential equations). We find no mention of independence in the text and we do not believe that the dependence impacts the relevance of this model to state estimation.

- *Page 5, lines 26-27: Kalman filter equation is normally solved by in lower space with the covariance represented by ensembles, there is no need for explicit representation.*

The reviewer makes a nice point that we have now included in the manuscript:

One remedy is to solve the Kalman filter equation in reduced space with the covariance represented by ensembles rather than being explicitly represented. Instead, we design a whole-domain method that is computationally efficient and provides a good first guess for the boundary control problem.

- *Page 14, equations (A1), (A2). Omega is the time derivative of theta. They are not completely independent.*

We agree with the reviewer's statement, but admit that we are unclear as to the larger relevance of this statement.

- *Page 25, the 2nd to the last sentence in Figure 7 caption is clear not a correct sentence.*

We have revised the sentence to the following.

The local minimum in the first two dimensions is no longer an extremum in the other two dimensions or the combined three-dimensional space.

References

- Baker, G. L., and J. P. Gollub (1990), *Chaotic Dynamics: An Introduction*, Cambridge University Press.
- Bennett, A. F. (1992), *Inverse methods in physical oceanography*, 346 pp., Cambridge Monographs, Cambridge University Press.
- Bennett, A. F. (2002), *Inverse Modeling of the Ocean and Atmosphere*, 234 pp., Cambridge University Press.
- Bonekamp, H., G. J. Van Oldenborgh, and G. Burgers (2001), Variational assimilation of tropical atmosphere-ocean and expendable bathythermograph data in the Hamburg Ocean Primitive Equation ocean general circulation model, adjusting the surface fluxes in the tropical ocean, *JGR*, 106(C8), 16,693–16,709.
- Casti, J. L. (1982), Recent developments and future perspectives in nonlinear system theory, *SIAM Review*, 24(3), 301–331.
- Casti, J. L. (1985), *Nonlinear system theory*, vol. 175, Academic Press.
- Cohn, S. E., and D. P. Dee (1988), Observability of discretized partial differential equations, *SIAM Journal on Numerical Analysis*, 25(3), 586–617.
- Courtier, P., J.-N. Thepaut, and A. Hollingsworth (1994), A strategy for operational implementation of 4D-Var, using an incremental approach, *Quart. J. Roy. Meteor. Soc.*, 120, 1367–1387.
- Ferron, B., and J. Marotzke (2003), Impact of 4D-variational assimilation of WOCE hydrography on the meridional circulation of the Indian Ocean, *Deep Sea Research Part II: Topical Studies in Oceanography*, 50, 2005–2021.
- Fukumori, I., and P. Malanotte-Rizzoli (1995), An approximate Kalman Filter for ocean data assimilation: an example with an idealized Gulf-Stream model, *J. Geophys. Res.*, 100(C4), 6777–6793.
- Fukumori, I., J. Benveniste, C. Wunsch, and D. B. Haidvogel (1993), Assimilation of sea surface topography into an ocean circulation model using a steady-state smoother, *J. Phys. Oceanogr.*, 23, 1831–1855.
- Hall, M. C. G., D. G. Cacuci, and M. E. Schlesinger (1982), Sensitivity analysis of a radiative-convective model by the adjoint method, *J. Atmos. Sci.*, 39, 2038–2050.
- LeDimet, F., and O. Talagrand (1986), Variational algorithm for analysis and assimilation of meteorological observations: Theoretical aspects, *Tellus*, 38A, 97–110.
- Malanotte-Rizzoli, P., and E. Tziperman (1995), The oceanographic data assimilation problem: Overview, motivation and purposes, in *Modern approaches to data assimilation in ocean modeling*, edited by P. Malanotte-Rizzoli, Elsevier Oceanography Series.
- Marchal, O. (2014), On the observability of oceanic gyres, *Journal of Physical Oceanography*, 44(9), 2498–2523.
- Stammer, D. (1997), Global characteristics of ocean variability estimated from regional TOPEX/Poseidon altimeter measurements, *J. Phys. Oceanogr.*, 27, 1743–1769.
- Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C. N. Hill, and J. Marshall (2002), The global ocean circulation during 1992-1997, estimated from ocean observations and a general circulation model, *J. Geophys. Res.*, 107 (C9), 3118.
- Stammer, D., K. Ueyoshi, A. Kohl, W. G. Large, S. A. Josey, and C. Wunsch (2004), Estimating air-sea fluxes of heat, freshwater, and momentum through global ocean data assimilation, *Journal of Geophysical Research*, 109(C5).
- Thacker, W. C., and R. B. Long (1988), Fitting dynamics to data, *J. Geophys. Res.*, 93, 1227–1240.
- Tziperman, E., and W. C. Thacker (1989), An optimal-control/adjoint-equations approach to studying the oceanic general circulation, *J. Phys. Oceanogr.*, 19, 1471–1485.
- Tziperman, E., L. Stone, M. A. Cane, and H. Jarosh (1994), El Niño chaos: overlapping of resonances between the seasonal cycle and the Pacific ocean-atmosphere oscillator, *Science*, 264(5155), 72–74.
- Wunsch, C. (2010), *Discrete Inverse and State Estimation Problems. With Geophysical Fluid Applications*.
- Wunsch, C., and P. Heimbach (2007), Practical global oceanic state estimation, *Physica D-nonlinear Phe-*

nomena, 230((1-2)), 192–208.

Controllability, not chaos, key criterion for ocean state estimation

Geoffrey Gebbie¹ and Tsung-Lin Hsieh^{1,2,3}

¹Department of Physical Oceanography, Woods Hole Oceanographic Institution, Woods Hole, MA, USA

²Summer Student Fellow, Woods Hole Oceanographic Institution, Woods Hole, MA, USA

³Princeton University, Princeton, NJ, USA

Correspondence to: Geoffrey Gebbie (ggebbie@whoi.edu)

Abstract.

The Lagrange multiplier method for combining observations and models (i.e., the adjoint method or “4D-VAR”) has been avoided or approximated when the numerical model is highly nonlinear or chaotic. This approach has been adopted primarily due to difficulties in the initialization of low-dimensional chaotic models, where the search for optimal initial conditions by gradient descent ~~algorithm~~ algorithms is hampered by multiple local minima. Although initialization is an important task for numerical weather prediction, ocean state estimation usually demands an additional task – solution of the time-dependent surface boundary conditions that result from atmosphere-ocean interaction. Here, we apply the Lagrange multiplier method to an analogous boundary control problem, tracking the trajectory of the forced chaotic pendulum. Contrary to previous assertions, it is demonstrated that the Lagrange multiplier method can track multiple chaotic transitions through time, so long as the boundary conditions render the system controllable. Thus, the nonlinear timescale poses no limit to the time interval for successful Lagrange multiplier-based estimation. That the key criterion is controllability, not a pure measure of dynamical stability or chaos, illustrates the similarities between the Lagrange multiplier method and other state estimation methods. The results with the chaotic pendulum suggest that there is no fundamental obstacle to ocean state estimation with eddy-resolving, highly-nonlinear models, especially when using an improved first-guess trajectory.

15 1 Introduction

The most complicated, and probably most realistic, numerical models of the ocean circulation are eddy-resolving ocean general circulation models (e.g., Arbic et al., 2010; Maltrud et al., 2010; Griffies et al., 2015). Such models are a natural choice in ocean state estimation, the combination of models and observations to reconstruct our best estimate of what the ocean has actually done (e.g., Stammer et al., 2002a). Here, we restrict our focus to state estimation as the transient reconstruction of the ocean state over a previous finite-interval of time where observations are available, following the convention of Wunsch et al. (2009). In order to unambiguously diagnose physical mechanisms of interest, the ocean state must be dynamically consistent: a solution to the dynamical equations of motion without any unphysical terms. The Lagrange multiplier method (e.g., Thacker and Long, 1988; Wunsch, 2010), sometimes called the adjoint method (e.g., Hall et al., 1982; Tziperman and Thacker, 1989), “4D-VAR” (e.g., Courtier et al., 1994; Ferron and Marotzke, 2003), or variational data assimilation (e.g., LeDimet and Talagrand, 1986;

Bonekamp et al., 2001; Bennett, 2002), is a method that satisfies both of these criteria, unlike the Kalman filter (e.g., Fukumori and Malanotte-Rizzoli, 1995) or nudging techniques (e.g., Malanotte-Rizzoli and Tziperman, 1995).

For the Lagrange multiplier method to be successful in state-of-the-art ocean models, two major issues need to be addressed: (1) the high dimensionality of the forward model and estimation problem, and (2) the nonlinearity of ocean models at increasingly fine resolution. Research conducted by the ECCO (Estimating the Circulation and Climate of the Ocean) Consortium (Stammer et al., 2002b, 2004) has demonstrated that (1), the dimensionality of many million state variables ~~is not a fundamental problem~~ presents a challenge, but, insofar as a solution can be found that fits the ocean data, it can be overcome and it is does not pose a fundamental obstacle. One caveat is that the convergence of the optimization process may be slower than hoped, but this is primarily an issue of computational efficiency. Regarding nonlinearity (2), the adjoint model has the same stability characteristics as the forward model, as the eigenvalues of linearized state transition matrix are the same as the transpose of the matrix (Palmer, 1996). Therefore, nonlinearity in the forward model may be accompanied by an unstable adjoint model and Lagrange multipliers that grow exponentially with time. When the Lagrange multiplier method is used to enforce a nonlinear constraint such as a chaotic model, the search for a solution becomes iterative and the Lagrange multipliers provide gradient information that is used to minimize an objective function that describes the model fit to observations (e.g., Marotzke et al., 1999). For a bounded objective function with growing gradients, multiple local minima are present that complicate the search for a global minimum (e.g., McShane, 1989). Even sophisticated gradient descent algorithms such as the variable-storage quasi-Newton method (Nocedal, 1980; Gilbert and Lemaréchal, 1989) can become stalled in a local minimum and are not guaranteed to fit the observations adequately. For example, Lea et al. (2000) used the Lorenz (1963) model to conclude that the "adjoint does not tend to useful sensitivity values," echoing previous concerns with simple, chaotic models (e.g., Gauthier, 1992; Miller et al., 1994a; Tanguay et al., 1995).

Due in part to the concerns raised about nonlinearity in simple models, the method of Lagrange multipliers has rarely been applied to realistic models over time windows longer than the eddy scale. For example, some studies restricted the time windows to be short enough that unstable modes would not grow too large (e.g., Schröter et al., 1993; Cong et al., 1998). The Southern Ocean State Estimate was produced with an approximate version of the method of Lagrange multipliers, where the Lagrange multipliers are calculated by an adjoint model with artificially large diffusivities that stabilize the model (Mazloff et al., 2010). Such an approach is not guaranteed to work, as the Lagrange multipliers of the stabilized model have no simple relation with those of the eddy-resolving model. The iterative search technique could then be led in the opposite direction as the true solution, as was shown to occur in a quasi-geostrophic ocean model (Köhl and Willebrand, 2002). We are aware of only one case where the unmodified method of Lagrange multipliers was applied to an eddy-permitting ocean GCM over a timescale longer than the eddy scale of a few months (Gebbie et al., 2006). Contrary to expectation given by the simple chaotic models, an acceptable fit was found to oceanographic observations over a one-year interval in the Northeast Atlantic Ocean (Gebbie, 2007). No clear explanation for these disparate results has been put forward.

In this research, we wish to re-examine (2), the influence of nonlinear models on the method of Lagrange multipliers and ocean state estimation. Is the adjoint method useless with a highly-nonlinear or chaotic system, as studies with low-dimensional chaotic models suggest? Here we posit that the initialization problem that has informed much of the current thinking about

the Lagrange multiplier method is not the relevant analogy for ocean state estimation. ~~The As has been documented in detail by many authors including the textbook of Bennett (1992), the~~ ocean state estimation problem ~~may be~~ is better described as a time-variable boundary value problem because synoptic atmospheric variability acts as an external forcing on the ocean (Section 2). Given our relatively uncertain knowledge regarding air-sea fluxes, the ocean state estimation ~~problem should be~~ is rightfully considered a time-variable boundary value problem where both the initial conditions and boundary conditions must be found. ~~In this case~~ For example, Bennett (2002) described an estimation method for the external forcing, initial and boundary conditions that solves the Euler-Lagrange equations for a linear model. In addition, ocean state estimation with general circulation models (e.g., Köhl and Stammer, 2008) takes the surface forcing to be part of the control vector. Because the effect of nonlinearity is seen as the major roadblock for application of the Lagrange multiplier method, we isolate this effect by choosing a model that is highly nonlinear but low-dimensional: the forced, chaotic pendulum (Section 2). Toy models are worth revisiting because the dynamics are comparatively simple to understand, the nonlinear coupling to periodic forcing has been shown to be important in atmosphere-ocean dynamics (e.g., Tziperman et al., 1994), and these models have strongly influenced when the Lagrange multiplier method has been deployed to realistic ocean problems. We will show that previous toy models have sometimes been misinterpreted.

15 Rather than developing a new state-of-the-art data assimilation technique, we proceed by taking the existing Lagrange multiplier method and developing diagnostics regarding when and why it succeeds or fails, as evaluated by the ability to fit observations. Relative to the initialization problem, the prospects for a successful state estimate are shown to be improved in the boundary control problem, even if one uses a highly nonlinear model such as the forced, chaotic pendulum (Section 3). If the chaotic nature of the model is not a roadblock, what is the relevant criterion for success with the Lagrange multiplier

20 method? Our results with the chaotic pendulum suggest that “controllability,” defined as the ability to move from one arbitrary state to another by control adjustments, is the relevant parameter. The implications of these results is that there is a wide variety of situations where the Lagrange multipliers of an ocean general circulation model (GCM) are useful, and that previous GCM results can be explained in this context (Section 4).

2 Lagrange multiplier method

25 2.1 Pendulum model and synthetic data

The fixed, single pendulum can be modeled as a nonlinear or linear set of equations, and it can also be easily modified to be stable or unstable. In many ways, the pendulum is a more flexible and easily interpreted physical system than the often-used Lorenz (1963) equations that approximate atmospheric convection. The relevance of the pendulum to the ocean is obviously indirect, but much of the community’s knowledge of state estimation has been formed ~~by~~ through the intuition of ~~studies of~~ simple models. The motion of the forced pendulum is described by the deterministic equation (Baker and Gollub, 1990),

$$\frac{d^2\theta}{dt^2} + \frac{1}{q} \frac{d\theta}{dt} + \frac{g}{l} \sin \theta = f(t), \quad (1)$$

where θ is the displacement angle from vertical, q is a damping coefficient, g is gravitational acceleration, l is the pendulum length, and $f(t) = b \cos(\omega_d t)$ is an external forcing term. Later, the external forcing will be broken into a first guess and a perturbation, $f(t) = f_0 + \delta f(t)$, where the first-guess is set to periodic forcing, $f_0(t) = b \cos(\omega_d t)$. With parameters $q = 100$ s, $g/l = 1.0 \text{ s}^{-2}$, $b = 1.5 \text{ rad s}^{-2}$, and $\omega_d = 2/3 \text{ s}^{-1}$, the pendulum is chaotic (here defined as extreme sensitivity to initial conditions). Following the numerical implementation in Appendix A, the state vector is defined, $\mathbf{x}(t) = [\omega(t) \theta(t)]^T$, $\mathbf{x}(t) = [\omega(t) \theta(t)]^T$, where T is the vector transpose. Matrices and vectors are indicated in boldface. The state has dimension $M = 2$ and the forcing vector has dimension 1. The evolution of the state is succinctly written,

$$\mathbf{x}(t + \Delta t) = \mathcal{L}[\mathbf{x}(t), f(t)], \quad (2)$$

where the model state is stepped from time t to $t + \Delta t$, and \mathcal{L} is the discretized, nonlinear operator that represents equation (1).

10 In the ocean model case, the state would correspond to velocities and property fields, and the external forcing would include air-sea momentum, heat, and freshwater fluxes.

We consider an ‘‘identical twin’’ experiment where the true solution is known (solid line, Figure 1), and we observe the pendulum angle episodically through time with normally-distributed random errors of standard deviation, $\sigma_\theta = 0.5$ rad. In most oceanographically-relevant cases, observations have already been collected over some fixed time interval ($0 \leq t \leq T$).

15 Here, observations, $y(t)$, are taken at a set of N_y evenly-spaced times with an time interval of $\Delta t_y = T/(N_y - 1)$.

2.2 Cost function

We proceed by defining a least-squares cost function to be minimized. The data-based contribution to the cost function, J_d , measures the squared misfit between the model and observations:

$$J_d = \frac{1}{N_y} \sum_{i=0}^{N_y-1} \frac{[\theta(i\Delta t_y) - y(i\Delta t_y)]^2}{\sigma_\theta^2} \quad (3)$$

20 where the penalty is weighted by the number of observations, N_y , and their standard error, σ_θ , such that the expected value of J_d is near 1. As we have imposed Gaussian error statistics, minimizing this least-squares cost function also leads to the maximum likelihood solution (e.g., Jazwinski, 1970). In matrix-vector notation, equation (3) becomes,

$$J_d = \sum_{i=0}^{N_y-1} \frac{N_y-1}{N_y} [\mathbf{E}\mathbf{x}(i\Delta t_y) - y(i\Delta t_y)]^T W^{-1} [\mathbf{E}\mathbf{x}(i\Delta t_y) - y(i\Delta t_y)], \quad (4)$$

where $y(t)$ is a scalar, \mathbf{E} is the observational matrix that samples the observable part of the state and has dimension $1 \times M$, and W is a weight. Comparison of the first term in (4) to equation (3) shows that $W = N_y \sigma_\theta^2$. While it is unconventional to transpose the scalar data-model misfit in equation (4), we retain this notation so that the equations are applicable to cases where multiple observations are available at each time.

A second contribution to the cost function includes two terms that constrain the difference between our posterior and prior estimates of the initial conditions and forcing,

$$30 \quad J_0 = [\mathbf{x}(0) - \mathbf{x}_0(0)]^T \mathbf{S}_x^{-1} [\mathbf{x}(0) - \mathbf{x}_0(0)] + \sum_{t=0}^{N_t-1} [f(i\Delta t) - f_0(i\Delta t)]^T S_f^{-1} [f(i\Delta t) - f_0(i\Delta t)] \quad (5)$$

where $\mathbf{x}_0(0)$ is the first-guess initial conditions, there are N_t model timesteps, $f_0(t)$ is the first-guess forcing, and S_x and S_f are weights that restrict the size of the perturbations to 5rad and 10rad s⁻², respectively. Here we seek values of $\mathbf{x}(t)$ and $f(t)$ that minimize the sum, $J' = J_d + J_0$, but the stationary point found by individually minimizing the values $dJ'/d\mathbf{x}(t)$ and $dJ'/df(t)$ will almost certainly violate the model constraint in equation (2). We enforce the model constraint by appending a

5 Lagrange multiplier term to the combined cost function,

$$J = J_d + J_0 - 2 \sum_{t=0}^{N_t-1} \mu(i\Delta t + \Delta t)^T \{ \mathbf{x}(i\Delta t + \Delta t) - \mathcal{L}[\mathbf{x}(i\Delta t), f(i\Delta t)] \}, \quad (6)$$

where $\mu(t)$ is a Lagrange multiplier, and the scaling with “2” is helpful in later derivations and does not change the numerical value of J because the quantity inside curly brackets vanishes. Now the cost function can be minimized by independently setting the partial derivatives of J with respect to the state, the forcing, and the Lagrange multipliers to zero. This problem will

10 be solved using a gradient-descent method (detailed later) that is excellent at finding the nearest minimum. If the first guess is good, then the closest minimum may actually be the global minimum (e.g., Pires et al., 1996), and therefore we design an improved first guess next.

2.3 First-guess trajectory

Minimizing J requires a first-guess of the full model trajectory, $\mathbf{x}_0(t)$. A sensible and common approach is to use the observa-

15 tion at initial time, $y(0)$, to inform the initial conditions for the state, $\mathbf{x}_0(0)$. Then, the first-guess forcing, $f_0(t) = b \cos(\omega t)$, is used to drive the model forward in time. In this case, the state at any time, τ , can be computed directly from the initial state,

$$\mathbf{x}_0(\tau) = \mathcal{L}_{K-1}[\dots[\mathcal{L}_1[\mathcal{L}_0[\mathbf{x}_0(0), f_0(0)], f_0(\Delta t)] \dots], f_0(K\Delta t - \Delta t)] = \mathcal{R}(\tau, 0)[\mathbf{x}_0(0)], \quad (7)$$

where \mathcal{L}_k indicates the nonlinear model operator at timestep k , $K = \tau/\Delta t$ is the number of timesteps between $t = 0$ and $t = \tau$, and the state transition matrix, $\mathcal{R}(m, n)$, defines the aggregate, nonlinear model step to time m from n . In the following, we

20 refer to this trajectory as the “standard” first-guess state.

For a nonlinear system, and a chaotic system in particular, this first-guess trajectory usually diverges from the already-

collected observations at some point, and thus can be ruled out as a possible solution *a priori*. When the pendulum initial conditions are imperfectly known, the range of possible pendulum trajectories expands greatly with time, even if the forcing evolution is perfectly known (background shading, Figure 1). Normally-distributed initial perturbations to the truth with stan-

25 dard deviation of 1 rad s⁻¹ in angular velocity and 0.5 rad in the initial angle lead to a divergence of roughly 200 rad between extreme trajectories (background shading, Figure 1). The angle is not renormalized when the angle is greater or less than π , and thus the angle records a history of how many times the pendulum has rotated. If no information about the initial angular velocity is available, a reasonable assumption is that $\omega = 0$ with some large error, but the pendulum trajectory with this initial velocity and the correct initial angle diverges from truth in less than 5 seconds (first dashed line, Figure 1). In the case where

30 the initial velocity is known perfectly but the initial angle is observed with an initial error of 0.5 rad (second dashed line, Figure 1), the trajectory follows truth for 30 seconds before eventually diverging. As the time interval of interest increases, any uncertainty in the initial conditions will ultimately lead to a divergence between truth and this first-guess model trajectory.

While these sample model trajectories may seem overly naive, the first guess trajectory used for ocean state estimation usually has similar characteristics: usage of an observation at the initial time, some prior knowledge of the forcing, and a freely-running forward model.

2.4 An improved first guess

- 5 The aforementioned standard approach does not use the observational information already in hand that could inform the time evolution of the forcing. There are many methods that are available to update the forcing, such as the Kalman filter (e.g., Keppenne et al., 2005), but these methods rival or exceed the method of Lagrange multipliers in computational cost because of the explicit representation of the solution covariance matrix (Fukumori, 2002). ~~Here~~ [One remedy is to solve the Kalman filter equation in reduced space with the covariance represented by ensembles rather than being explicitly represented. Instead,](#)
- 10 we design a [whole-domain](#) method that is computationally efficient and provides a good first guess for the boundary control problem.

Here we seek an update to the initial conditions and the forcing (i.e., $\omega_1(0) = \omega_0(0) + \delta\omega$, $\theta_1(0) = \theta_0(0) + \delta\theta$, $f_1(t) = f_0(t) + \delta f(t)$), that takes the observations into account. For small perturbations, we derive a linearized equation for the change to the state at the time of the first observation, $t = \Delta t_y$,

$$\mathbf{x}_1(\Delta t_y) = \mathbf{x}_0(\Delta t_y) + [\Pi_{i=1}^K \mathbf{A}((K-i)\Delta t) \mid \Pi_{i=1}^{K-1} \mathbf{A}((K-i)\Delta t)\mathbf{B} \mid \Pi_{i=1}^{K-2} \mathbf{A}((K-i)\Delta t)\mathbf{B} \mid \dots \mid \mathbf{B}] \begin{pmatrix} \delta\omega(0) \\ \delta\theta(0) \\ \delta f(0) \\ \delta f(\Delta t) \\ \vdots \\ \delta f(\Delta t_y - \Delta t) \end{pmatrix} + \epsilon \tag{8}$$

15

where \mathbf{x}_1 is the ‘‘improved’’ first-guess, $K = \Delta t_y / \Delta t$ is the number of model time steps from $t = 0$ to $t = t_y$, $\mathbf{A}(t) = \partial \mathcal{L} / \partial \mathbf{x}(t)$ is the tangent-linear model, $\mathbf{B} = \partial \mathcal{L} / \partial f(t)$ is constant in time, and ϵ is the error due to linearization. We define the column vector of perturbations in equation (8) to be the control vector, \mathbf{u} , so that the equation becomes:

$$\mathbf{x}_1(\Delta t_y) = \mathbf{x}_0(\Delta t_y) + \mathbf{C}\mathbf{u} + \epsilon, \tag{9}$$

- 20 where \mathbf{C} is the controllability (or reachability) matrix (e.g., Dahleh and Diaz-Bobillo, 1999; Wunsch, 2010).

The observation, $y(\Delta t_y)$, and the combination of equations (7) and (9) provides one constraint,

$$y(\Delta t_y) = \mathbf{E}\mathcal{R}(\Delta t_y, 0)[\mathbf{x}_0(0)] + \mathbf{E}\mathbf{C}\mathbf{u} + \mathbf{n} \tag{10}$$

where the controllability matrix can be calculated given the trajectory, $\mathbf{x}_0(t)$, and \mathbf{n} is the misfit. Here we minimize the squared misfit,

$$J_1 = \{y(\Delta t_y) - \mathbf{E}\mathcal{R}(\Delta t_y, 0)[\mathbf{x}_0(0)] - \mathbf{E}\mathbf{C}\mathbf{u}\}^T W^{-1} \{y(\Delta t_y) - \mathbf{E}\mathcal{R}(\Delta t_y, 0)[\mathbf{x}_0(0)] - \mathbf{E}\mathbf{C}\mathbf{u}\} + \mathbf{u}^T \mathbf{Q}^{-1} \mathbf{u}, \tag{11}$$

where \mathbf{Q} is a block diagonal matrix with \mathbf{S}_x and S_f on the diagonal. We solve with the method of total inversion (Tarantola and Valette, 1982),

$$5 \quad \mathbf{u} = \mathbf{W}\mathbf{C}^T\mathbf{E}^T[\mathbf{E}\mathbf{C}\mathbf{W}\mathbf{C}^T\mathbf{E}^T + \mathbf{Q}]^{-1}\{y(\Delta t_y) - \mathbf{E}\mathcal{R}(\Delta t_y, 0)[\mathbf{x}_0(0)]\}, \quad (12)$$

where we update the state transition and controllability matrices iteratively. The full nonlinear model is run with the updated controls to produce the improved first-guess trajectory, $\mathbf{x}_1(t)$. The algorithm proceeds sequentially $N_y - 1$ times, where the terminal state from one segment becomes the initial condition for the next.

2.5 Solution for Lagrange multipliers

10 We obtain the sensitivity of J to the initial conditions by taking the partial derivative,

$$\frac{\partial J}{\partial \mathbf{x}(0)} = 2\mu(\Delta t) + 2\mathbf{S}_x^{-1}[\mathbf{x}_1(0) - \mathbf{x}_0(0)] + 2\mathbf{E}^T\mathbf{W}^{-1}[\mathbf{E}\mathbf{x}_1(0) - y(0)], \quad (13)$$

where the improved first guess, $\mathbf{x}_1(t)$, is used. Taking the derivative with respect to the other set of unknowns, we find,

$$\frac{\partial J}{\partial f(t)} = 2\mathbf{B}^T\mu(t + \Delta t) + 2S_f^{-1}[f_1(t) - f_0(t)]. \quad (14)$$

With knowledge of these gradients, we could improve the initial conditions and forcing, but both equations (13) and (14) depend upon the Lagrange multipliers, $\mu(t)$, that we must solve for first.

Extending the partial derivative of J with respect to $\mathbf{x}(t)$ and $\mu(t)$ at all times, we recover the Euler-Lagrange (or “normal”) equations. The Lagrange multipliers are determined by timestepping backward in time:

$$\mu(t) = \mathbf{A}(t)^T\mu(t + \Delta t) + \mathbf{E}^T\mathbf{W}^{-1}[\mathbf{E}\mathbf{x}_1(t) - y(t)], \quad (15)$$

where the last term on the right hand side only appears if an observation, $y(t)$, is available. Initial conditions are given by,

$$20 \quad \mu(T) = \mathbf{E}^T\mathbf{W}^{-1}[\mathbf{E}\mathbf{x}_1(T) - y(T)]. \quad (16)$$

Equations (15) and (16) are collectively known as the *adjoint model* (e.g., Bugnion et al., 2006), where the model-observation misfit is part of the adjoint model forcing, $\mathbf{E}^T\mathbf{W}^{-1}[\mathbf{E}\mathbf{x}_1(t) - y(t)]$. Now the result of equation (15) can be substituted into equations (13) and (14) to solve for the gradients.

In summary, the method of Lagrange multipliers is implemented with the following steps. Starting with a guess for the initial conditions and forcing, $\mathbf{x}_0(0)$ and $f_0(t)$, we improve the first guess and solve for the full trajectory, $\mathbf{x}_1(t)$, using the method of total inversion. The adjoint model is then run backward in time to solve for the sensitivity of J with respect to the two types of unknowns, $\mathbf{x}(0)$ and $f(t)$. Here we use a quasi-Newton gradient descent optimization (Nocedal, 1980) to update these uncertain control parameters. Because the model is nonlinear, the tangent-linear model, $\mathbf{A}(t)$, will depend upon the nonlinear model trajectory, and we re-run the full nonlinear model to get an updated trajectory that will replace $\mathbf{x}_1(t)$. Forward-adjoint model integrations are repeated until J has an acceptable value by a χ^2 statistical test.

3 Results

3.1 Tracking chaotic transitions

5 Synthetic observations of the pendulum angle are generated every 2.5 seconds over a 50 second time interval, where a random error of 0.5 rad is added to every observation. We first illustrate the futility of a brute force search for the optimal initial conditions by re-running the forward model with combinations of the initial angle and angular velocity in the neighborhood of the truth. The data contribution to the cost function (equation 3) is then evaluated for each forward model trajectory, giving rise to a complex topology where the global minimum is not immediately visible (Figure 2). That the topology of the cost
10 function is not conducive for gradient descent search was previously documented in models of convection, quasi-geostrophic flow, and the oceanic double gyre model (e.g., Miller et al., 1994b; Köhl and Willebrand, 2003; Lea et al., 2006). When the initial angle and angular velocity are slightly perturbed from the truth, the cost function values can become extremely large due to the divergence of trajectories. Furthermore, the cost function varies irregularly with many local extrema at locations other than the true solution. The basin of attraction of the true solution, defined in analogy to a drainage basin on a topographic map,
15 is much smaller than the observational uncertainty, and thus, it is likely that the iterations of the adjoint method will converge to a local, non-global, minimum.

We start the application of the methods of Section 2 by determining the first-guess trajectory. Here we implement a model timestep of 0.01 s over a 50s integration time and thus the control vector has 5000 forcing variables and 2 initial condition variables. The first-guess initial conditions, forcing, and trajectory are calculated according to Section 2.3. Despite the assump-
20 tion of linearity in the calculation of the first guess, the improved first guess has a trajectory that is nearly consistent with the error bars of the observations (gray line, top panel, Figure 3). The first guess also tracks the rapid transitions in the interval, $12s < t < 20s$, where revolutions of the pendulum occur due to chaotic dynamics. These results contrast with a seemingly reasonable first-guess trajectory that is determined by a model simulation initialized with the first observation ($\theta(0) = y(0)$) and zero angular velocity that goes off track in less than 5 seconds (see “standard first guess”, top panel, Figure 3). Similar
25 first-guesses are common in ocean models, where the circulation field is started from rest with the assumption that geostrophic balance will equilibrate the velocity field rapidly. Our more sophisticated, but still linear, method of deriving an improved first guess makes the Lagrange multiplier method more likely to succeed.

The improved first-guess trajectory is a better fit to the data in large part due to the updated initial angular velocity (middle panel, Figure 3). Starting with an angular velocity of about 2 rad/s, this trajectory has frequent changes in the sign of angular
30 velocity consistent with reversals in pendulum rotation. Conversely, the standard trajectory has a long period of strictly positive angular velocity ($10 s < t < 40 s$) that is inconsistent with the observations. Another important factor is the position of the pendulum around 10 seconds after the start of the integration, when small differences in the state become greatly magnified. The true pendulum trajectory then enters a period where several revolutions occur successively. The inaccurate initial velocity causes errors at this critical time of instability and thus the trajectories diverge.

For θ and ω , the difference between the improved-first-guess and final (Lagrange multiplier method) trajectories is smaller than the changes brought about by the first-guess improvement itself. The method of Lagrange multipliers acts similarly to a

combined filter-smoother that simultaneously takes into account past and future observations, leading to an angular velocity evolution with somewhat smaller range while still fitting the observations. Consequently, the evolution of the pendulum angle is also smoother, with fewer variations at the observational sampling frequency of $1/2.5$ s. The full impact of the Lagrange multiplier method only becomes clear when considering the external forcing in the following section.

3.2 Reconstruction of the forcing

The improved first-guess trajectory better fits the observations than a standard first-guess, but there are tradeoffs in the estimated forcing (bottom panel, Figure 3). In order to track the chaotic transitions of the pendulum, the improved first-guess trajectory makes forcing adjustments that are sometimes strong and abrupt in order to compensate for previous errors in the trajectory. In other words, these adjustments take the forcing evolution farther away from the truth that was used in the standard first-guess trajectory. This is reminiscent of the small-scale features that are added to the surface forcing of ocean models in order to fit observations (e.g., Stammer et al., 2002b), although our case is not a compensation for inaccurate model dynamics because we are operating under a perfect model assumption. This tradeoff is probably unacceptable for those wishing to physically interpret the forcing field, and indicates that the improved first-guess estimate is not a good final solution despite fitting the observations. The power of the Lagrange multiplier method is now clear; not only is the final estimate smoother than the first guess, the final forcing estimate accurately reproduces the amplitude and frequency of the true forcing: $f(t) = 1.5 \cos(2t/3)$.

Our identical twin experiment permits a comparison with the truth to diagnose actual errors even at times without observations. While the improved first guess appeared to fit the data well, the misfit to the truth displays considerable structure, including a large deviation around $t = 45$ s to values that are inconsistent with the observations (top panel, Figure 4). Such a deviation reflects an inaccurate interpolation between data points during the construction of the improved first guess trajectory. The first guess also appears to overfit the observations, as this estimate deviates from the truth in the neighborhood of observations with large error (e.g., $t = 32.5, 42.5$ s). The final estimate, on the other hand, hovers near 0 for the entire time interval, with a standard deviation of 0.46 rad, very near to the actual observational uncertainty of 0.50 rad. The final estimate reproduces 72% of the variance in the observational error, computed by comparison of the estimated to true observational error. Visually, the final estimate is closer to the truth than the observations over the majority of the time interval, indicating that the Lagrange multiplier method filters out the observational noise even in this chaotic system.

For the angular velocity and forcing (middle and bottom panels, Figure 4), the Lagrange multiplier method reproduces the truth despite the imperfect, sparse observations and the chaotic model dynamics. The suppression of the abrupt and large changes in forcing is not simply a smoothing or averaging of the forcing, but instead is seen to reflect the true forcing, as evidenced by the deviation from true forcing being small. The method of Lagrange multipliers is therefore superior to the first-guess estimate because all components of the solution, both the state and forcing, can be physically interpreted.

3.3 Influence of the data stream

In the case where only two observations are available ($\theta(0) = -2 \pm 0.5, \theta(50) = 10 \pm 0.5$), the time between observations is greater than both the fundamental frequency and the nonlinear timescale of the pendulum. The nonlinear timescale is here

defined as the time interval that the tangent linear model well-approximates the nonlinear dynamics, which depends upon the size of initial perturbation. Despite the long time between observations, the estimated trajectory fits both observations via the Lagrange multiplier method (top left panel, Figure 5). Thus, there appears to be no lower limit on the number of observations
5 necessary in order to produce an acceptable state estimate with this model. Of the two steps in our method, it is the first-guess calculation that is responsible for fitting the data within their errors, and the optimization with Lagrange multipliers doesn't substantially change or improve this estimate. When comparing the estimate to the true trajectory that was withheld from the reconstruction method, differences larger than 10 rad exist (bottom left panel, Figure 5). Thus, reconstruction of the full, partially-unobserved trajectory without any intervening observations is a challenging task, as expected. Here, we emphasize
10 that the first goal in state estimation is to find any model trajectory that fits the observations, and that in realistic cases we will not know whether the model interpolates between the observations in the correct way. We recognize that a chaotic model usually has many trajectories that satisfy the initial and final times, and thus, any one trajectory is unlikely to reconstruct the truth at all intervening times.

In the case that many ($N_y = 20$) observations are taken but the standard error is large (5 rad), all observations are again fit
15 within their 1σ error bars. Strictly speaking, the data appear to be overfit, as 32% of the points are expected to reside outside the one standard error level but none do. This fit has a lower standard deviation, 3.3 rad, than that expected by the observational error of 5.0 rad, suggesting that the numerical model is adding information that is complementary to the observations. Only in short time intervals does the estimate differ from the truth by more than 5 rad, such as near the observation of the 2σ outlier at $t = 25$ s. Unlike the case where 20 observations were taken with a smaller standard error of 0.5 rad in the previous section, this
20 estimate is only partially successful at filtering noise out of the observations. The correlation coefficient between the estimated and actual observational error is $r = 0.43$, indicating that some noise remains. Both case studies in this section indicate that neither the quality nor the quantity of the data stream affect whether the Lagrange multiplier method can be successful with a chaotic model.

3.4 Influence of the number of controls

25 The previous section addresses cases where the forcing is adjusted at every timestep, leading to 5002 control variables. After application of the Lagrange multiplier method, the resulting value of the data-based component of the cost function (equation 3) depends upon both the number of control parameters and the number of observations (Figure 6). The value of J_d is always below 1 when 5002 control variables are defined, consistent with the previously-reported results. Here we test the null hypothesis that the model is consistent with the observations, and we perform a χ^2 statistical test with N_y degrees of freedom as is appropriate for independent observations. Our one-sided test statistic is the value of J_d where 5% of cases are expected to have larger values by chance. For 5002 controls, we find that all values are small enough that the null hypothesis cannot be rejected at the 5% insignificance level (area above black line, Figure 6). Thus, the Lagrange multiplier method is expected
5 to acceptably fit the observations if enough controls are available. In the ocean state estimation problem, all air-sea fluxes are uncertain and temporally variable so there is no shortage of controls that can be defined.

We investigate the effect of a decrease in the number of controls by redefining the external forcing control perturbation. For N_u forcing controls, we define,

$$f(t) = f_0(t) + \Gamma(t) \begin{pmatrix} \delta f(0) \\ \delta f(T/(N_u - 1)) \\ \delta f(2T/(N_u - 1)) \\ \vdots \\ \delta f(T) \end{pmatrix}, \quad (17)$$

10 where $\Gamma(t)$ is a row vector that performs linear interpolation in time, and $\delta f(t)$ is only defined at N_u control times. This formulation enforces some temporal correlation in the external forcing. Alternatively, this could be accomplished using a nondiagonal weighting matrix, \mathbf{S}_f .

For a given number of observations, a decrease in the number of controls leads to a decrease in the likelihood of a successful fit to the data. The initialization problem is equivalent to the case with two boundary-control variables, and Figure 6 suggests that
 15 the Lagrange multiplier method will not produce a good fit to data, as documented by previous works. The criterion of a good fit also depends upon the number of observations, where more observations decrease the likelihood of success. To understand why the data can or cannot be fit, consider that each observation gives a constraint of the type documented in equation (10). If all of these constraints are enforced simultaneously, the problem is formally underdetermined when the number of controls exceeds the number of observations, and it is generally likely that a solution exists. The simple interpretation that the number
 20 of controls must exceed the number of observations does not strictly hold due to the logarithmic scale in Figure 6. Even when the problem is formally underdetermined, a singular value decomposition analysis of the controllability matrix, \mathbf{C} , reveals that not all controls are independent and that the data cannot always be fit perfectly. Here, we identify formally underdetermined cases with a J_d value that is unacceptably high (Figure 6), thus the solvability condition is sometimes violated. Such a result indicates that the controllability matrix has an effective rank less than the number of observations, showing the importance of
 25 this quantity as a diagnostic measure of the conditioning of the estimation problem. In practice, the singular values need not be strictly zero, as a large discrepancy between the magnitudes of singular values can give ill-conditioning.

We also find cases where the gradient descent method is capable of navigating the complex cost function topology with Lagrange multiplier sensitivity information. A slice of J along $\omega(0)$ and $\theta(0)$ is focused on a region of phase space that appears to contain a local minimum, although this is not the true solution (left panel, Figure 7). In the initial control problem, the
 30 optimization would proceed in these two dimensions and be trapped by the local minimum. Taking a 2D slice of the cost function in the dimension of the initial angular velocity and forcing, $\delta f(0)$, however, the same location may no longer be a local minimum in the expanded phase space (right panel, Figure 7). In our example, the cost function can be further minimized by decreasing $\delta f(0)$, and the optimization process may eventually get out of the trap in $\omega(0)/\theta(0)$ space. Thus, additional dimensions in the optimization space can sometimes alleviate problems with the gradient descent algorithm.

3.5 Influence of prior forcing information

The previous examples in Section 3 proceed with prior information that the forcing is periodic with an accurate magnitude and phase. A good analogy is the regular forcing of solar insolation on the ocean surface. Here, we test the performance of the Lagrange multiplier method with inaccurate prior information about the forcing, as is a more realistic analogy to the uncertainty of air-sea fluxes. In particular, our first guess of the forcing, $f_0(t)$, is systematically biased by decreasing b from 1.5 to 0.75 rad s^{-2} . The trajectory driven by inaccurate forcing is no worse than the previous cases with accurate forcing due to the dominance of the chaotic dynamics of system (Figure 8). Using the same observations as shown in Figure 3, we find that the chaotic pendulum trajectory is tracked over multiple nonlinear timescales despite this more stringent test. In this case, however, the forcing estimate still contains errors relative to the true forcing calculated with $b = 1.5 \text{ rad s}^2$, and some high-frequency structures remain in $f(t)$ (see “improved first guess” in bottom panel, Figure 8). If instead the Lagrange multiplier method is started from the standard first guess, a smoother and more accurate estimate of the forcing is obtained at the expense of not fitting the data as well (see “final estimate” in bottom panel). Any remaining irregular structures can be handled by imposing temporal correlations as was done in Section 3.4. If such measures are not taken, the investigator must take care to decide what elements of the forcing represent true variability and which are compensating for model error. In our simple system of equations, model errors and forcing errors are mathematically equivalent. In state estimates with eddy-resolving GCMs, however, smallscale forcing variability is found near oceanic fronts and the investigator must determine on a case-by-case basis to what extent it reflects real variability.

4 Discussion

4.1 Relation to Kalman Filter/Smoother

Our suggestion that controllability is a key criterion brings our understanding of the Lagrange multiplier method into closer consistency with the Kalman filter and smoother. Both methods solve the same least-squares problem, and the solution of a linear problem should not depend upon the chosen method (e.g., Wunsch, 2010). Fukumori et al. (1993) found that the problem must be controllable for the Kalman filter/smoothing to be successful. In addition, the chaotic Lorenz (1963) model was tracked with the Kalman filter/smoothing over time windows much longer than the nonlinear timescale when the system was completely controllable (i.e., all estimated quantities are uncertain and are treated as control variables) (Evensen, 1997). Our results suggest that the equivalence of the Kalman filter/smoothing and Lagrange multiplier method may be extended to nonlinear problems, thus explaining why the chaotic estimation problem may be solved by the Lagrange multiplier method.

To recover the true trajectory of a system, observability is also important (Fukumori et al., 1993; Marchal, 2014). Wunsch (1996) states, as the estimation problem is the dual of the control problem (Fukumori et al., 1993; Marchal, 2014). For the linear problem, Cohn and Dee (1988) showed that completely observability implies asymptotic stability of the Kalman filter/smoothing. Defining observability and controllability conditions for nonlinear state estimation problems is difficult Casti (1985). In practice, the important criterion is ability to solve equation (10). Strictly speaking, the solution criteria will therefore depend upon both the controllability matrix, C , and the observational matrix, E , which combines the issues of observability and controllability. Here,

we suggest the operational definition that a system is effectively controllable when the solution to (10), generalized to multiple observations, exists.

5 Related to the idea of observability, Wunsch (1996) stated that, “Problems owing to the multiple minima in the cost function can always be overcome by having enough observations to keep the estimates close to the true state.” To evaluate this state-
ment, we emphasize that there are two levels of successful reconstruction: 1) one that accurately fits the data, and 2) one that
accurately fits the truth at all times and locations. Criterion 1) has been our metric for success in this work, as in real-world
problems, criterion 2) cannot be tested. Here we have shown that only controllability is necessary for 1), even with a nonlinear
system. In addition, we show that the data can still be fit even if very few observations are available, as an off-track estimate can
be righted by precise adjustments to the forcing (recall Figure 3). That short-lived forcing adjustments can put the estimate on
10 track is likely a consequence of the nonlinear dynamics of our particular problem, although we believe that an eddy-resolving
ocean general circulation model could behave the same way. Our interpretation is consistent with work in the control of chaotic
systems. Engineers have described the control of a chaotic system as being "easier" than control of other systems because the
necessary control adjustments are very small (e.g., Ott et al., 1990).

4.2 Comparison of controllability and stability metrics

15 In the section, we compare criteria for the success of the Lagrange multiplier method. Previously-suggested criteria include
Lyapunov exponents or other stability metrics of the tangent-linear model (e.g., Lea et al., 2000). The tangent linear matrix
has eigenvalues with absolute value greater than one when linearized about a state in the upper-half plane ($\pi/2 < \theta < 3\pi/2$),
reflecting the divergence of neighboring nonlinear trajectories when a pendulum perturbed towards the horizontal is more
rapidly accelerated downwards. Conversely, the lower-half plane is locally linearly stable. The unforced pendulum with ini-
20 tial conditions in the upper-half plane is episodically-unstable, until damping brings the pendulum permanently into a stable
configuration in the lower half-plane.

Here we investigate the influence of stability versus that of nonlinearity. The pendulum is a useful system because it is easily
modified to have 4 distinct dynamical states: 1) nonlinear, unstable, 2) nonlinear, stable, 3) linear, stable, and 4) linear, unstable.
Case 1) is the original dynamical equation for the pendulum (equation 1). By restricting the phase space to the lower-half plane,
25 the pendulum is locally linearly stable at all times, although it is still nonlinear (Case 2). When the pendulum is linearized by
the small-angle approximation with the linear term of the Taylor series expansion ($\sin \theta \approx \theta$, see Appendix A2), we obtain a
linear, stable model (Case 3). If instead the pendulum is linearized around its apex, the sign of θ in the linearized equation is
reversed, rendering the system linear but unstable (Case 4).

We revisit the problem of estimating the initial angle when θ is observed at the final time. A synthetic observation is generated
30 by running the model with initial displacement of $-\pi/6$ rad and zero velocity. Assuming a perfect model and observation, the
shape of the cost function is generated by changing the initial conditions and evaluating J . In the two nonlinear cases, a slice of
the cost function contains many local minima that emerge when the time window is extended from 5 to 50 seconds (Figure 9,
lower panels). The cost function in the nonlinear, stable case deviates from a parabola because the state transition matrix is
non-self-adjoint and non-normal growth occurs (Farrell 1989; Farrell and Moore 1993). Thus, even nonlinear models that are

stable are subject to local minima, and linear stability is not always a good metric to determine whether a gradient descent search will successfully find the global minimum.

Conversely, the linear, unstable case does yield a parabolic cost function (Figure 9, *upper panels*), implying that instability
5 does not impede the search for the minimum. Again, local linear stability doesn't appear to be a good metric for determining the presence of local minima, because an unstable system may yield a well-behaved function. This example reinforces the counterintuitive relationship between stability and local minima, where a linearly stable system does not have a paraboloidal cost function but an unstable system does. While this reversed relationship does not always hold, linear stability metrics are not reliable. We suggest that controllability is a better metric, but note that controllability and stability are not unrelated, as a
10 system with a growing unstable mode could lead to a controllability matrix that effectively drops rank.

4.3 Relevance to ocean state estimation

In the Introduction, we remarked on the only ocean state estimate known to the authors that successfully implemented the Lagrange multiplier method in an eddy-permitting ocean GCM without any modification to the adjoint model (Gebbie et al., 2006). In light of the results of this work, a combination of factors appears to have been responsible for that success. While
15 the ocean model had $1/6^\circ \times 1/6^\circ$ horizontal resolution and contained mesoscale eddies (Figure 10), the resolution was not adequate to fully resolve the eddies. Also, the model domain of the Northeast Atlantic Ocean was a relatively quiescent one. Both factors likely led to the ocean model being more linear than other studies. In addition, the adjoint model of a coarse-resolution twin was used to form an improved first guess, which would improve the likelihood of success with gradient descent much as our method did here. Perhaps most importantly, the ocean state estimate included air-sea control fields that were
20 updated every 10 days, leading to a total of 5.5×10^6 control variables. Given the rapid adjustment of the ocean to barotropic waves, it is likely that the system passed the controllability criterion derived in this work. Controllability could be numerically evaluated in a GCM by a series of impulse functions: a dynamical equivalent to the passive response recorded by transit time distributions (e.g., Delhez et al., 1999; Haine and Hall, 2002). Open questions include whether the deep ocean is completely controllable by surface boundary conditions, and whether ocean data require variability at timescales shorter than 10 days to
25 be introduced through the surface forcing.

5 Conclusions

There is no fundamental obstacle to constraining a highly nonlinear model to observations using the Lagrange multiplier method. On the basis of research primarily with toy models, chaotic systems were thought to represent such an obstacle if the estimation time window was too long. Here we find that the trajectories of the nonlinear pendulum can be tracked over
30 multiple rapid transitions that are due to chaotic dynamics. The Lagrange multiplier method is successful under the condition that enough boundary controls are available through time, and that the system passes a test of controllability. In the case of the pendulum, the rank of the controllability matrix is a better metric to predict a success of state estimation rather than a measure

of dynamical stability. The ocean state estimation problem is analogous to the problem posed here; uncertain air-sea fluxes contain large errors that require control adjustments through time.

Our implementation of the Lagrange multiplier method includes a step to construct a good first guess that helps the iterative gradient descent search. The first-guess method has been developed with implementation in an ocean GCM in mind. Specifically, sub-problems are defined over the interval between observations and thus require less memory than a whole-domain approach. In addition, we suggest that the particular first-guess method of this work is not the only way to produce a good first guess, and that other methods would bring the first-guess state close enough to the truth to increase the likelihood of success. A good example is the Green's function method (e.g., Stammer and Wunsch, 1996; Menemenlis et al., 2004) that selects a subset of the full control variables and makes some linearity assumptions. Following up the Green's function optimization with a gradient descent search with the Lagrange multiplier method is therefore a worthwhile research goal. The results of this work suggest that ocean state estimation should continue with the Lagrange multiplier method and models that resolve higher and higher resolution physics.

Appendix A: Numerical implementation of pendulum

A1 Nonlinear pendulum

The forced, nonlinear pendulum is governed by the following equation,

$$\frac{d^2\theta}{dt^2} + \frac{1}{q} \frac{d\theta}{dt} + \frac{g}{l} \sin \theta = f(t), \quad (\text{A1})$$

where the symbols were defined in Section 2. Discretizing in time, we obtain the symbolic form of the model equation used in the main text:

$$\frac{d\mathbf{x}(t)}{dt} = \frac{d}{dt} \begin{pmatrix} \omega(t) \\ \theta(t) \end{pmatrix} = \begin{pmatrix} -\frac{1}{q} \omega(t) - \frac{g}{l} \sin \theta(t) + f(t) \\ \omega(t) \end{pmatrix}. \quad (\text{A2})$$

If the system is discretized with a forward Euler timestep of time Δt , the discrete-time state space realization is:

$$\begin{pmatrix} \omega(t+1) \\ \theta(t+1) \end{pmatrix} = \begin{pmatrix} (1 - \Delta t/q) \omega(t) - (g \sin \theta(t)/l + f(t)) \Delta t \\ \Delta t \omega(t) + \theta(t) \end{pmatrix}, \quad (\text{A3})$$

or simply,

$$\mathbf{x}(t+1) = \mathcal{L}[\mathbf{x}(t), f(t)], \quad (\text{A4})$$

where \mathcal{L} is a nonlinear operator due to the sine function. For use in the Euler-Lagrange equations, we also produce the following linearized operators:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{x}(t)} \equiv \mathbf{A}(t) = \begin{pmatrix} 1 - \Delta t/q & -g \cos \theta(t) \Delta t/l \\ \Delta t & 1 \end{pmatrix}, \quad \frac{\partial \mathcal{L}}{\partial f(t)} \equiv \mathbf{B} = \begin{pmatrix} \Delta t \\ 0 \end{pmatrix}. \quad (\text{A5})$$

Here we use a second-order Taylor timestepping (i.e., midpoint forward Euler method) for increased accuracy. We code the
 5 tangent-linear model in accordance with differentiation rules for numerical codes (Griewank, 2000), and we run this linearized
 model with perturbations to all elements of $\delta\mathbf{x}(t)$ and $\delta f(t)$ to recover the values. Numerical parameters include: $\Delta t = 0.01\text{s}$,
 $\omega_{true}(0) = 1.2959\text{rad/s}$, $\theta_{true}(0) = -2.4667\text{rad}$, and the forcing phase, $\phi_{true}(0) = 0.3412\text{rad}$.

A2 Linear, stable pendulum

The linear, stable pendulum is derived with the small-angle approximation. This approximation is a linearization around zero
 5 displacement

$$\begin{pmatrix} \omega(t + \Delta t) \\ \theta(t + \Delta t) \end{pmatrix} = \begin{pmatrix} 1 - g\Delta t & -g\Delta t/l \\ \Delta t & 1 \end{pmatrix} \begin{pmatrix} \omega(t) \\ \theta(t) \end{pmatrix}. \quad (\text{A6})$$

The more general tangent linear model is re-linearized around a changing nonlinear model trajectory.

Acknowledgements. We thank Geir Evensen, Armin Köhl, Olivier Marchal, and Eli Tziperman for discussions on this topic over the last
 decade, and to Jacques Verron for his note that has encouraged this work. GG also acknowledges Carl Wunsch ~~for his guidance on originally~~
 10 ~~selecting~~, [Patrick Heimbach, Detlef Stammer, and Julio Sheinbaum for guidance getting started on](#) this project. GG was funded through
 the Ocean and Climate Change Institute of the Woods Hole Oceanographic Institution. TLH was funded by the Woods Hole Oceanographic
 Institution Summer Student Fellowship program through the U.S. National Science Foundation.

References

- 15 Arbic, B., Wallcraft, A., and Metzger, E.: Concurrent simulation of the eddy general circulation and tides in a global ocean model, *Ocean Modelling*, 32, 175–187, 2010.
- Baker, G. L. and Gollub, J. P.: *Chaotic Dynamics: An Introduction*, Cambridge University Press, 1990.
- Bennett, A. F.: *Inverse methods in physical oceanography*, Cambridge Monographs, Cambridge University Press, 1992.
- Bennett, A. F.: *Inverse Modeling of the Ocean and Atmosphere*, Cambridge University Press, 2002.
- 20 Bonekamp, H., Van oldenborgh, G. J., and Burgers, G.: Variational Assimilation of Tropical Atmosphere–Ocean and expendable bathythermograph data in the Hamburg Ocean Primitive Equation ocean general circulation model, adjusting the surface fluxes in the tropical ocean, *JGR*, 106, 16 693–16 709, 2001.
- Bugnion, V., Hill, C., and Stone, P. H.: An adjoint analysis of the meridional overturning circulation in a hybrid coupled model, *Journal of climate*, 19, 3751–3767, 2006.
- Casti, J. L.: *Nonlinear system theory*, vol. 175, Academic Press, 1985.
- 25 Cohn, S. E. and Dee, D. P.: Observability of discretized partial differential equations, *SIAM Journal on Numerical Analysis*, 25, 586–617, 1988.
- Cong, L. Z., Ikeda, M., and Hendry, R. M.: Variational assimilation of Geosat altimeter data into a two-layer quasi-geostrophic model over the Newfoundland ridge and basin, *Journal of Geophysical Research*, 103, 7719–7734, 1998.
- Courtier, P., Thepaut, J.-N., and Hollingsworth, A.: A Strategy for Operational Implementation of 4D-Var, Using an Incremental Approach, 30 *Quart. J. Roy. Meteor. Soc.*, 120, 1367–1387, 1994.
- Dahleh, M. A. and Diaz-Bobillo, I.: *Control of Uncertain Systems: A Linear Programming Approach*, Pergamon Press, 1999.
- Delhez, E., Campin, J., Hirst, A., and Deleersnijder, E.: Toward a general theory of the age in ocean modelling, *Ocean Modelling*, 1, 17–27, 1999.
- Evensen, G.: Advanced data assimilation for strongly nonlinear dynamics, *Monthly Weather Review*, 125, 1342–1354, 1997.
- 35 Ferron, B. and Marotzke, J.: Impact of 4D-Variational Assimilation of WOCE Hydrography on the Meridional Circulation of the Indian Ocean, *Deep Sea Research Part II: Topical Studies in Oceanography*, 50, 2005–2021, 2003.
- Fukumori, I.: A partitioned Kalman filter and smoother, *Monthly Weather Review*, 130, 1370–1383, 2002.
- Fukumori, I. and Malanotte-Rizzoli, P.: An approximate Kalman Filter for ocean data assimilation: an example with an idealized Gulf-Stream model, *J. Geophys. Res.*, 100, 6777–6793, 1995.
- Fukumori, I., Benveniste, J., Wunsch, C., and Haidvogel, D. B.: Assimilation of Sea Surface Topography into an Ocean Circulation Model Using a Steady-State Smoother, *J. Phys. Oceanogr.*, 23, 1831–1855, 1993.
- 5 Gauthier, P.: *Chaos and Quadri-dimensional Data Assimilation: A Study Based on the Lorenz Model*, *Tellus*, 44A, 2–17, 1992.
- Gebbie, G.: *Subduction in an Eddy-resolving State Estimate of the Northeast Atlantic Ocean*, Ph.D. thesis, Massachusetts Institute of Technology / Woods Hole Oceanographic Institution Joint Program in Oceanography, 2004.
- Gebbie, G.: Does eddy subduction matter in the Northeast Atlantic Ocean?, *J. Geophys. Res.*, 112, doi:10.1029/2006JC003 568, 2007.
- Gebbie, G., Heimbach, P., and Wunsch, C.: Strategies for Nested and Eddy-Permitting State Estimation, *J. Geophys. Res.*, 111, C10 073, 10 doi:10.1029/2005JC003094, 2006.
- Gilbert, J. C. and Lemaréchal, C.: Some Numerical Experiments with Variable-storage Quasi-Newton Algorithms, *Math. Program.*, 45, 407–435, 1989.

- Griewank, A.: *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, Society for Industrial and Applied Mathematics, Philadelphia, 2000.
- 15 Griffies, S. M., Winton, M., Anderson, W. G., Benson, R., Delworth, T. L., Dufour, C. O., Dunne, J. P., Goddard, P., Morrison, A. K., Rosati, A., et al.: Impacts on ocean heat from transient mesoscale eddies in a hierarchy of climate models, *Journal of Climate*, 28, 952–977, 2015.
- Haine, T. W. N. and Hall, T. M.: A generalized transport theory: Water-mass composition and age, *J. Phys. Oceanogr.*, 32, 1932–1946, 2002.
- Hall, M. C. G., Cacuci, D. G., and Schlesinger, M. E.: Sensitivity Analysis of a Radiative-Convective Model by the Adjoint Method, *J. Atmos. Sci.*, 39, 2038–2050, 1982.
- 20 Jazwinski, A.: *Stochastic Processes and Filtering Theory*, Academic Press, 1970.
- Keppenne, C., Rienecker, M., Kurkowski, N., Adamec, D., et al.: Ensemble Kalman filter assimilation of temperature and altimeter data with bias correction and application to seasonal prediction, *Nonlinear Processes in Geophysics*, 12, 491–503, 2005.
- Köhl, A. and Stammer, D.: Variability of the meridional overturning in the North Atlantic from the 50-year GECCO state estimation, *Journal of Physical Oceanography*, 38, 1913–1930, 2008.
- 25 Köhl, A. and Willebrand, J.: An Adjoint Method for the Assimilation of Statistical Characteristics into Eddy-resolving Ocean Models, *Tellus*, 54 (4), 406–425, 2002.
- Köhl, A. and Willebrand, J.: Variational Assimilation of SSH Variability from TOPEX/POSEIDON and ERS1 into an Eddy-permitting Model of the North Atlantic, *J. Geophys. Res.*, 108 (C3), 3092, 2003.
- Lea, D. J., Allen, M. R., and Haine, T. W. N.: Sensitivity Analysis of the Climate of a Chaotic System, *Tellus*, 52A, 523–532, 2000.
- 30 Lea, D. J., Haine, T. W. N., and Gasparovic, R. F.: Observability of the Irminger Sea Circulation using Variational Data Assimilation, 2006.
- LeDimet, F. and Talagrand, O.: Variational algorithm for analysis and assimilation of meteorological observations: Theoretical aspects, *Tellus*, 38A, 97–110, 1986.
- Lorenz, E. N.: Deterministic, Nonperiodic Flow, *J. Atmos. Sci.*, 20, 130–141, 1963.
- Malanotte-Rizzoli, P. and Tziperman, E.: The Oceanographic Data Assimilation Problem: Overview, Motivation and Purposes, in: *Modern approaches to data assimilation in ocean modeling*, edited by Malanotte-Rizzoli, P., Elsevier Oceanography Series, 1995.
- 35 Maltrud, M., Bryan, F., and Peacock, S.: Boundary impulse response functions in a century-long eddying global ocean simulation, *Environmental fluid mechanics*, 10, 275–295, 2010.
- Marchal, O.: On the observability of oceanic gyres, *Journal of Physical Oceanography*, 44, 2498–2523, 2014.
- Marotzke, J., Giering, R., Zhang, K. Q., Stammer, D., Hill, C., and Lee, T.: Construction of the Adjoint MIT Ocean General Circulation Model and Application to Atlantic Heat Transport Sensitivity, *J. Geophys. Res.*, 104, 529–547, 1999.
- Mazloff, M., Heimbach, P., and Wunsch, C.: An eddy-permitting Southern Ocean state estimate, *Journal of Physical Oceanography*, 40, 880–899, 2010.
- 5 McShane, E. J.: *The Calculus of Variations from the Beginning Through to Optimal Control Theory*, *SIAM J. Control Optim.*, 27, 916–939, 1989.
- Menemenlis, D., Fukumori, I., and Lee, T.: Using Green’s functions to calibrate an ocean general circulation model, *Monthly Weather Review*, 133, 1224–1240, 2004.
- Miller, R. N., Ghil, M., and Gauthiez, F.: Advanced Data Assimilation in Strongly Nonlinear Systems, *Journal of the Atmospheric Sciences*, 10 51, 1037–1056, 1994a.
- Miller, R. N., Zaron, E. D., and Bennett, A. F.: Data assimilation in models with convective adjustment, *Mon. Weath. Rev.*, 122, 2607–2613, 1994b.

- Nocedal, J.: Updating quasi-Newton matrices with limited storage, *Mathematics of computation*, 35, 773–782, 1980.
- Ott, E., Grebogi, C., and Yorke, J.: Controlling chaos, *Phys. Rev. Lett.*, 64, 1196–1199, 1990.
- 15 Palmer, T. N.: Predictability of the Atmosphere and Oceans: From Days to Decades, vol. 44 of *NATO ASI Series*, chap. 3, pp. 83–156, Springer, in *Decadal Climate Variability: Dynamics and Predictability*, 1996.
- Pires, C., Vautard, R., and Talagrand, O.: On extending the limits of variational assimilation in nonlinear chaotic systems, *Tellus A-Dynamic Meteorology and Oceanography*, 48, 96–121, 1996.
- Schröter, J., Seiler, U., and Wenzel, M.: Variational Assimilation of GEOSAT Data into an Eddy-resolving Model of the Gulf Stream
20 Extension Area, *J. Phys. Oceanogr.*, 23, 925–953, 1993.
- Stammer, D. and Wunsch, C.: The determination of the large-scale circulation of the Pacific Ocean from satellite altimetry using model Green’s functions, *Journal of Geophysical Research*, 101, 18 409–18 432, 1996.
- Stammer, D., Wunsch, C., Fukumori, I., and Marshall, J.: State Estimation in Modern Oceanographic Research, *EOS*, 83(27), 294–295, 2002a.
- 25 Stammer, D., Wunsch, C., Giering, R., Eckert, C., Heimbach, P., Marotzke, J., Adcroft, A., Hill, C. N., and Marshall, J.: The global ocean circulation during 1992-1997, estimated from ocean observations and a general circulation model, *J. Geophys. Res.*, 107 (C9), 3118, 2002b.
- Stammer, D., Ueyoshi, K., Kohl, A., Large, W. G., Josey, S. A., and Wunsch, C.: Estimating air-sea fluxes of heat, freshwater, and momentum through global ocean data assimilation, *Journal of Geophysical Research*, 109, 2004.
- Tanguay, M., Bartello, P., and Gauthier, P.: Four-dimensional Data Assimilation with a Wide Range of Scales, *Tellus*, 47A, 974–997, 1995.
- 565 Tarantola, A. and Valette, B.: Generalized Nonlinear Inverse Problems Solved using the Least Squares Criterion, *Reviews of Geophysics and Space Physics*, 20 (2), 219–232, 1982.
- Thacker, W. C. and Long, R. B.: Fitting dynamics to data, *J. Geophys. Res.*, 93, 1227–1240, 1988.
- Tziperman, E. and Thacker, W. C.: An Optimal-Control/Adjoint-Equations Approach to Studying the Oceanic General Circulation, *J. Phys. Oceanogr.*, 19, 1471–1485, 1989.
- 570 Tziperman, E., Stone, L., Cane, M. A., and Jarosh, H.: El Niño chaos: overlapping of resonances between the seasonal cycle and the Pacific ocean-atmosphere oscillator, *Science*, 264, 72–74, 1994.
- Wunsch, C.: *The Ocean Circulation Inverse Problem*, Cambridge University Press, New York, 1996.
- Wunsch, C.: *Discrete Inverse and State Estimation Problems. With Geophysical Fluid Applications*, 2010.
- Wunsch, C., Heimbach, P., Ponte, R., and Fukumori, I.: The global general circulation of the oceans estimated by the ECCO-Consortium,
575 *Oceanography*, 20, 88–103, 2009.




Figure 1. The rapid divergence of pendulum trajectories is indicated by the path density of trajectories (background shading), and the evolution of 3 sample trajectories: the “truth” or reference trajectory (*solid line*), a “first-guess” trajectory with incorrect initial angular velocity that diverges within 5s (*dashed line*), and a first-guess trajectory with incorrect initial angle, θ , that diverges after 30s (*other dashed line*). The path density of trajectories is computed with 10,000 forward integrations with normally-distributed perturbations about the truth (standard deviation: $\sigma_\omega = 1$ rad/s, $\sigma_\theta = 0.5$ rad).




Figure 2. Cost function values as a function of initial angle, θ , and angular velocity, ω , for an observational time window of 50s. The local minimum in this range of phase space found by the optimization (yellow dot) would is not be identifiable by eye. This figure was produced with the parameters, $T = 30$ s, $\Delta t_y = 1$ s, and $q = 4$ absolute minimum value of the cost function.

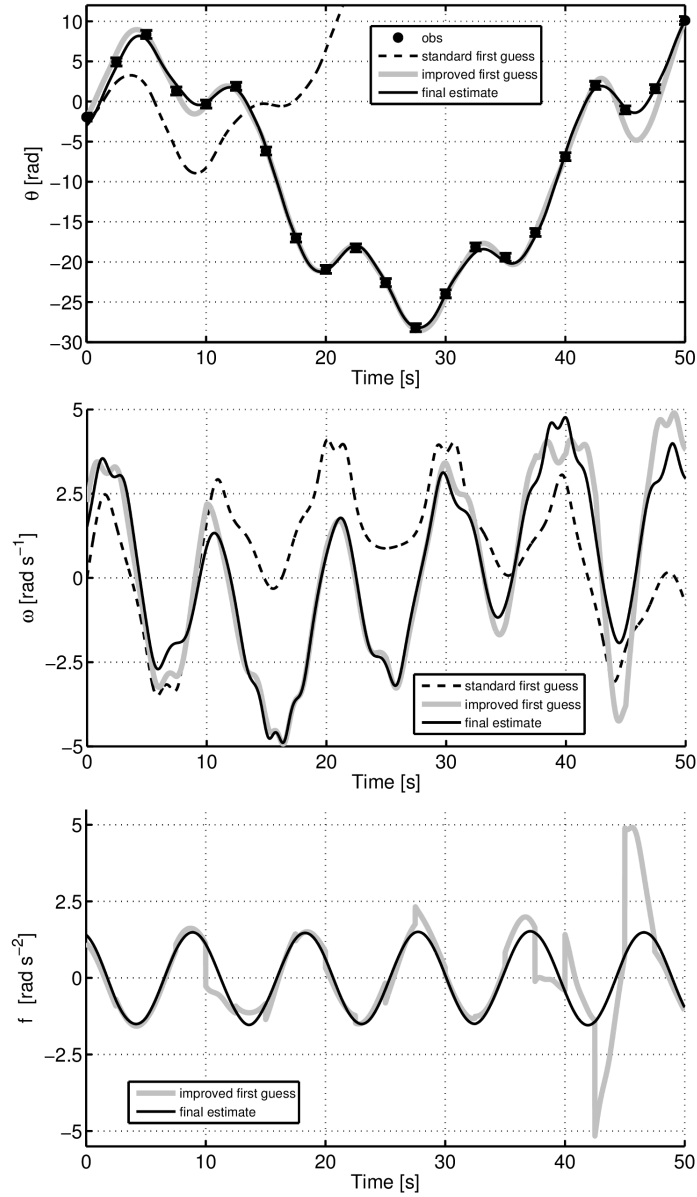


Figure 3. Control of the chaotic pendulum with an improved first guess and the Lagrange multiplier method. *Top panel:* Observations are taken every 2.5s with standard error of 0.5 rad (*circles with 1σ error bars*). The trajectory of the pendulum angle ($\theta(t)$, *top panel*), angular velocity ($\omega(t)$, *middle panel*), and the forcing ($f(t)$, *bottom panel*) are given for a standard first guess (*dashed line*), the improved first guess (*gray line*), and the final Lagrange multiplier-based estimate (*black line*). The standard first-guess forcing is not shown due to its similarity to the final estimate.

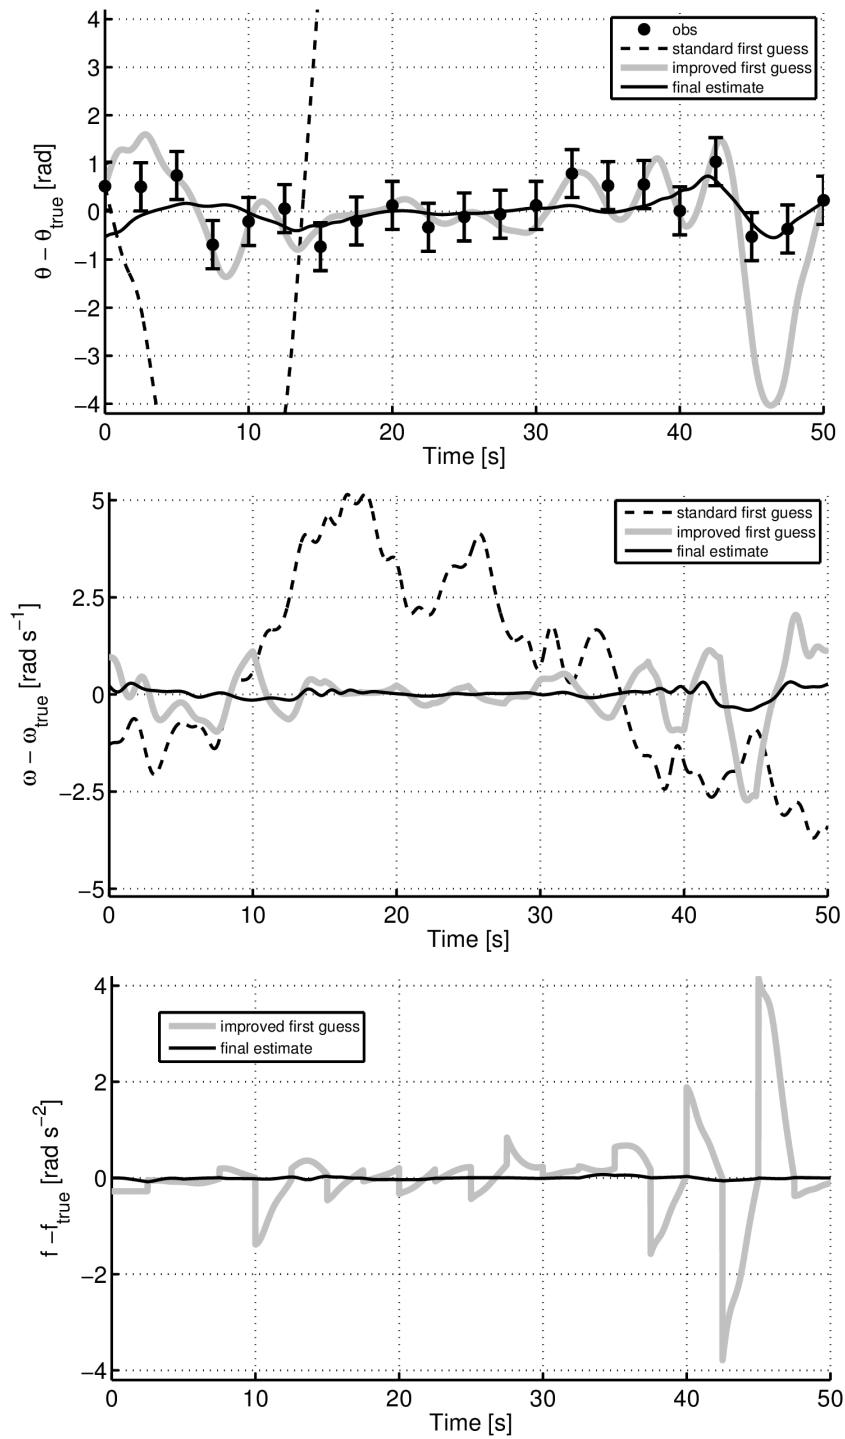


Figure 4. Comparison of the reconstructed pendulum and truth. *Top:* Difference between the observed (circles), standard first guess (dashed), improved first guess (solid gray line), and final estimate (solid, black line) of pendulum angle relative to the truth. The standard first-guess is off-scale for much of the panel. *Middle:* Similar but for angular velocity, ω . *Bottom:* Same but for forcing, f . The standard first-guess forcing is suppressed because it is identically zero.

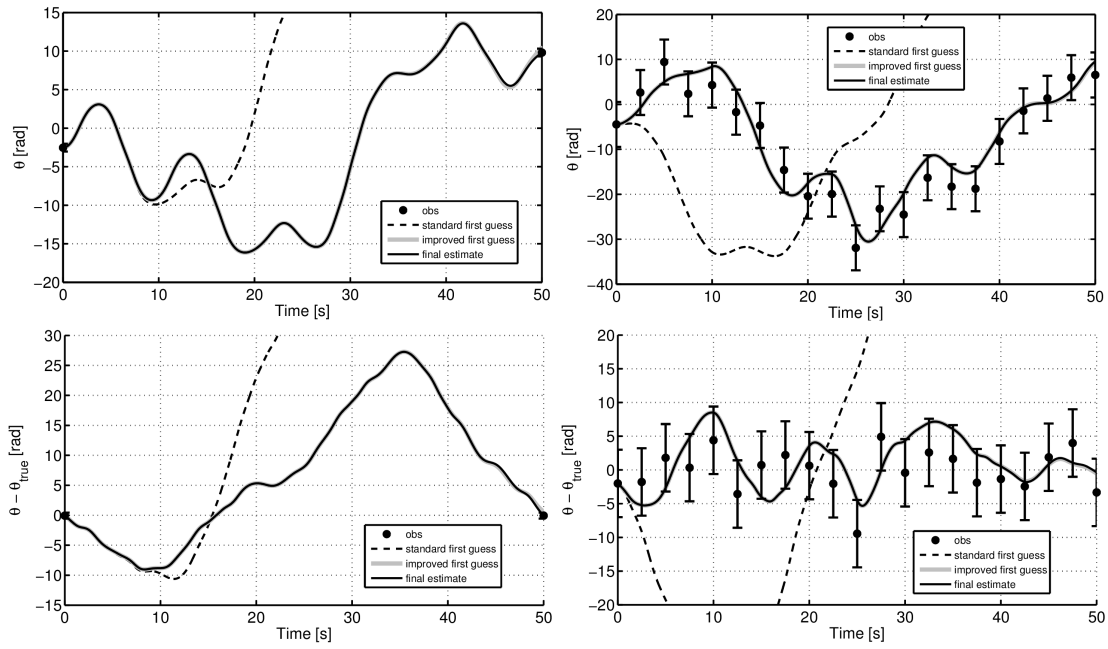


Figure 5. State estimation of the chaotic pendulum with a reduced set of 2 observations with standard error 0.5rad (*left column*) and a set of 20 observations with standard error 5rad. *Top:* Comparison of the observations (*circles with 1 σ error bars*), the standard first guess (*dashed*), the improved first guess (*gray solid line*), and the final state estimate (*solid black line*), as in Figure 3. *Bottom:* Similar to the top row, except all quantities are referenced to the truth, θ_{true} , as in Figure 4. Again the standard first guess is offscale for much of the time window. The improved first guess is nearly identical (and obscured) by the final estimate in all panels.




Figure 6. Influence of the number of observations and controls on the ability to track the chaotic pendulum. The base-10 logarithm of the data component of the cost function, $\log_{10}J_d$, is calculated as a function of the number of evenly-spaced observations over a 50 second window (*the abscissa*), and the number of effective degrees of freedom in the control perturbations to the forcing (*ordinate*). A χ^2 statistical test determines the limit where 95% of realizations are expected to have smaller J_d values (*thick, black line*); therefore, cases below this threshold represent an unacceptable fit to the data at the 5% insignificance level.




Figure 7. Escaping an apparent local minimum. *Left:* Cost function values as a function of $\theta(0)$ and $\omega(0)$, in a region of phas-phase space where a local minimum is present in this slice (yellow dot). The same cost function, but oriented along a slice with constant $\theta(0) = 0$ in the dimensions of $u_1 = \delta f(0) - \delta f(0)$ and $\omega(0)$. The local minimum is-in the first two dimensions is no longer an extremum in the other two dimensions or the combined three-dimensional space. This case used the parameters, $T = 10\text{s}$ and $\Delta t_y = 1\text{s}$, for illustration.

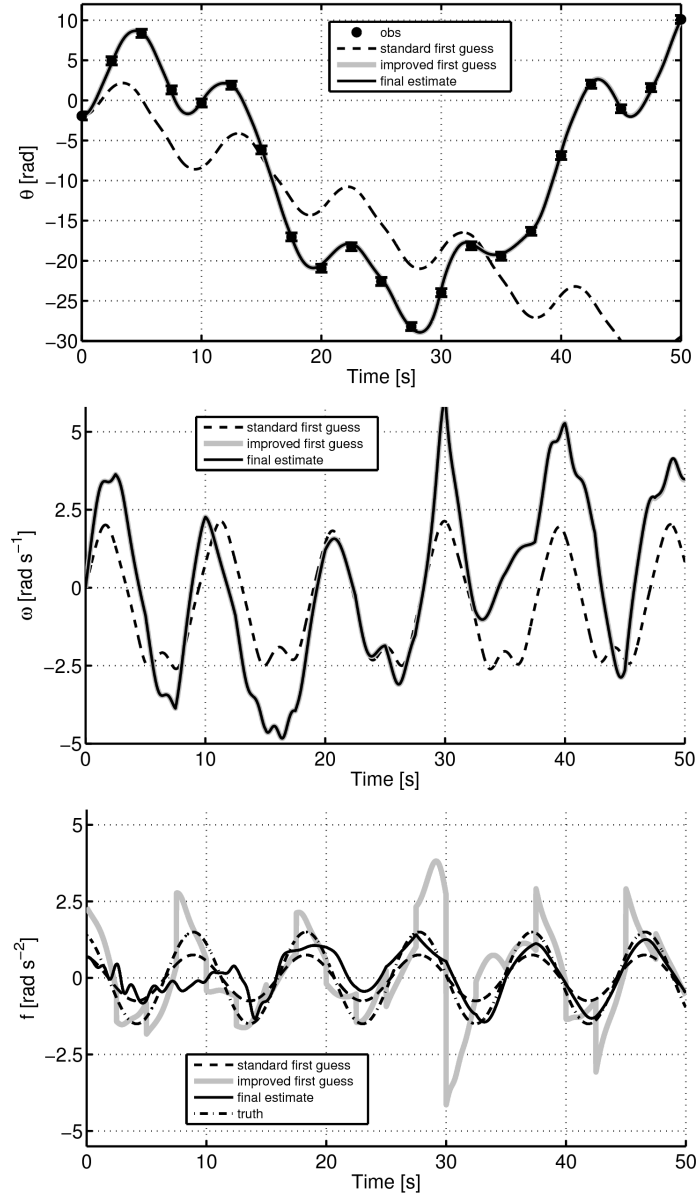


Figure 8. Control of the chaotic pendulum with an inaccurate first guess of the external forcing. *Top panel:* Observations are taken every 2.5s with standard error of 0.5 rad (circles with 1σ error bars). The trajectory of the pendulum angle ($\theta(t)$, *top panel*), angular velocity ($\omega(t)$, *middle panel*), and the forcing ($f(t)$, *bottom panel*) are given for a standard first guess (dashed line), the improved first guess (gray line), and the final Lagrange multiplier-based estimate (black line). In the forcing panel, we also include the true forcing (dash-dot line).

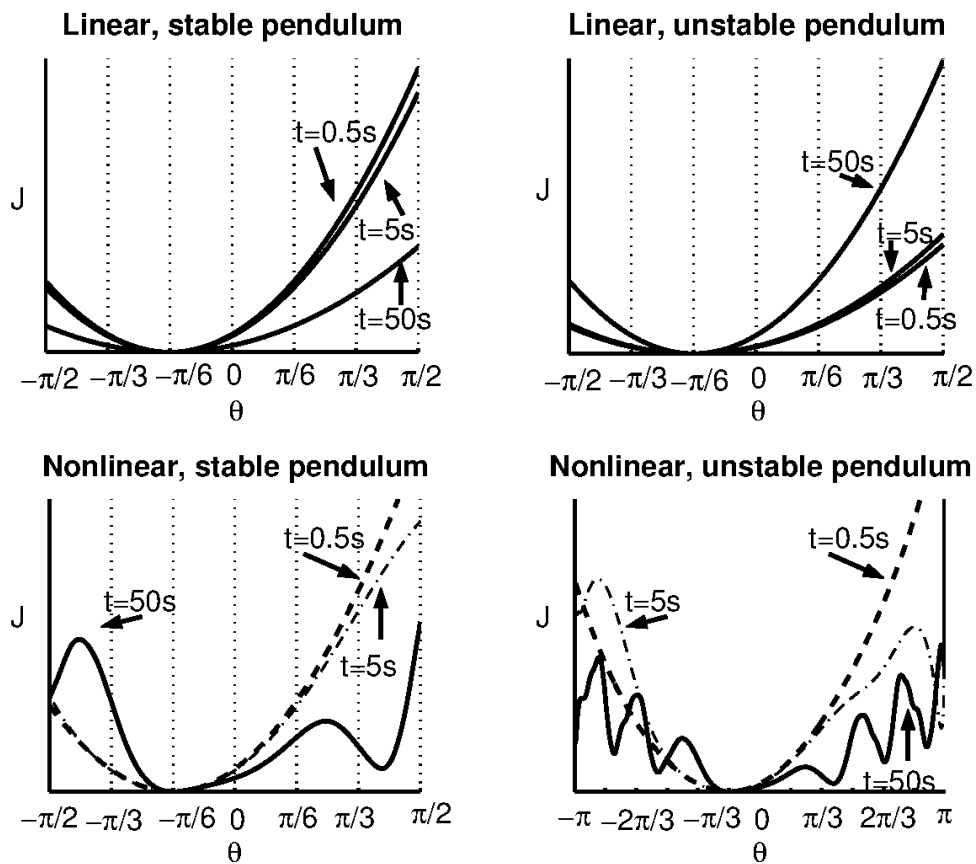


Figure 9. Cost function with respect to the initial pendulum angle. A synthetic observation was made from a model run with initial angle, $\theta = -\pi/6$. The time between the initial state and the cost function evaluation is 0.5, 5, or 50 seconds. *Upper left:* Linear, stable pendulum. *Upper right:* Linear, unstable pendulum. *Lower left:* Nonlinear, stable pendulum. *Lower right:* Nonlinear, unstable pendulum. Notice the wider scale for θ in the lower, right panel. The pendulum's dynamical regimes are further explained in the text. Reproduced with permission from Gebbie (2004).




Figure 10. Nested view of the $1/6^\circ$ regional state estimate of Gebbie et al. (2006) inside the 2° state estimate of Stammer et al. (2002b). Potential temperature at 310 meters depth, with a contour interval of 1°C , is shown. The boundary between the two estimates (*thick black line*) is discontinuous in temperature because of the open-boundary control adjustments. Reproduced with permission from Gebbie (2004).