

**Probabilistic State Estimation in Regimes of  
Nonlinear Error Growth**

by

W. Gregory Lawson

B.S. Geology & Geophysics  
Yale University, 1997

Submitted to the Department of Earth, Atmospheric, and Planetary Sciences  
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Atmospheric Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY


February 2005

© Massachusetts Institute of Technology 2005. All rights reserved.

Signature of Author.....  
Department of Earth, Atmospheric, and Planetary Sciences  
21 January 2005

Certified by.....  
James A. Hansen  
Assistant Professor of Meteorology

Accepted by.....  
Maria T. Zuber  
Chairman, Department Earth, Atmospheric, and Planetary Sciences  
E. A. Griswold Professor of Geophysics



ARCHIVES



# Probabilistic State Estimation in Regimes of Nonlinear Error Growth

by

W. Gregory Lawson

Submitted to the Department of Earth, Atmospheric, and Planetary Sciences  
on 21 January 2005, in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy in Atmospheric Science

## Abstract

State estimation, or data assimilation as it is often called, is a key component of numerical weather prediction (NWP). Nearly all implementable methods of state estimation suitable for NWP are forced to assume that errors remain in regimes of linear error growth and retain distributions of Gaussian uncertainty, yet nonlinear systems like the atmosphere can readily allow regimes of nonlinear error growth and, in turn, produce distributions of non-Gaussian uncertainty. State-of-the-art, ensemble-based methods of state estimation suitable for NWP are examined to gauge the consequences and relevance of violating the linear error growth assumption. For quite generic sources of non-Gaussian uncertainty, the methods are observed to fail, as they must, and the obtained analyses become probabilistically unreliable before becoming inaccurate.

The mis-positioning of coherent features is identified as a specific, geophysically relevant source of non-Gaussian uncertainty that can easily cause the state-of-the-art methods of state estimation to fail. However, an understanding of relevant phenomenology sometimes allows these same methods to remain successful owing to an available redefinition of the involved errors. The redefinition is phrased as an alternative error model. It is recognized and exploited that non-Gaussian additive Eulerian errors can come from Gaussian Lagrangian position errors. A two-step, augmented state vector approach is developed that is suitable for use with coherent features and that relies only on implementable methods of state estimation. By combining the dual Eulerian and Lagrangian state information into one vector, an ensemble can approximate their covariance, thus allowing each component's uncertainty to be reduced. The first step of the two-step approach reduces the feature position errors in an effort to render the residual additive errors Gaussian, thereby allowing the second step of an implementable state estimation method to proceed successfully. Philosophically, the two-step approach uses physical knowledge of the problem (as phrased by the error model) to compensate for neglected important non-Gaussian uncertainty structure in the state estimation process. The proposed two-step approach successfully allows use of implementable methods of state estimation to obtain probabilistically reliable analyses in regimes of nonlinear error growth, something unavailable using current standards.

Thesis Supervisor: James A. Hansen

Title: Assistant Professor of Meteorology



# Acknowledgments

*If you are going through hell, don't stop.* — Winston Churchill

Good advice, that. I nearly stopped many times along the way, but was able to persist only with the love, trust, and support of a large, strong “fan base.” Indeed, as they say, no man is an island; graduate school has a way of emphasizing this even as it tries to isolate its participants. So here in the acknowledgment section, I hope to specifically name many of the key participants in this truly group effort.

In my time here at MIT, I have had three different advisors, each of whom I thank for his own style of advising and support. Edmund Chang provided a good introduction to the ways that weather and graduate school work, Kerry Emanuel helped shepherd me into the realm of predictability and state estimation, and Jim Hansen has seen me through to the end, often with a lot of “above and beyond”-type support. Wearing many hats along the way, including those of teacher, agent, coach,  $\mathcal{F}$  operator, and friend, Super Jim is probably the single largest force behind my completion, and to him I owe a debt of gratitude not dissimilar to that I owe my parents, a debt that is most easily payed forward to future generations of graduate students, should I ever be in such a position to do so.

My thesis committee has been helpful along the journey through their advising and challenging of the material. Kerry and Paola Rizzoli have remained clear advocates of the work, and have each provided both valuable ideas and morale boosts. Carl Wunsch has also provided valuable ideas and has excellently played the (necessary) role of the critic. I thank Dennis McLaughlin for both his input on my thesis work and financial support through his awarded ITR grant (ACI-0121182). And I have had the pleasure of taking the T several times to Harvard Square to engage in enlightening conversations with Brian Farrell.

I have greatly appreciated the support found in my fellow classmates. I have enjoyed my membership in Jim’s research group, as well as many margaritas, with Vikram, Jon, and the two Andy’s (I don’t like to rank them 1 and 2). Jeff, Donnan, Rob, Nikki, and Mike, have been great friends as well as great bridge partners when time has allowed. I cherish my disc golf excursions with Jake as well as having a fellow student with whom to discuss state estimation. Peter has been an invaluable officemate, full of wit and insight. I have also greatly appreciated my interactions with Bridget, Juli, Markus, Baylor, Zan, Bill R., and Bill B. Some senior graduate students have also profoundly affected my experience here, most notably: Gerard, Rebecca, Nili, and Constantine.

Another MIT source of support is those that are employed to be supportive. Particular thanks goes to Kristen, Joel, Carol, Jackie, Roberta, Beth, Mary, and Vicki. Also, from

working late nights, I have had the pleasure of befriending many of the janitorial staff. I think that my relationship with the now-retired Tony Pallotta will be lifelong.

Beyond what I have already mentioned, I have dearly valued my interactions with Lodovica. As an instructor, she is a joy to have around, as well as a good coach for our forecasting team. I have had several useful discussions with Glenn Flierl about some subtleties in numerical modeling, and with Raffaele Ferrari about PDFs and turbulence. Also, I want to acknowledge the 11th hour adviser I have found in Dan Rothman. He is a great educator and of a very clear mind.

Outside of MIT, I have found Ross Hoffman to be an excellent source of feedback and advice. It has helped my work greatly to have one of its original pioneers so nearby and approachable. Similarly, I have valued my interactions with Craig Bishop; he is a bright young scientist with no shortage of energy or ideas.

I have greatly appreciated pursuing this degree in parallel with dear friends in other fields at other institutions. Ben, Greg, and Alex, all out-of-towners, have let me know that most of my problems and fears are closer to universal than idiosyncratic. Also, Johan and Trisha have been empathetic and supportive throughout. However, I have probably gained the most from my comrades in town: Doug, Gary, Ilan, and Liz have been ever-ardent supporters, and Matt, Jeff, and Andy have been brothers-in-arms. From getting drinks to hearing complaints, the throes and woes of a PhD require *ad hoc* group therapy, and they have always been there to supply it.

Before Jim, there was my undergraduate adviser Philip Bogden, who was very influential for me, right down to getting me to attend WHOI for a summer and to consider graduate school. Also, I have had my Uncle Blaine and Aunt Marie-Lousie as advisers my whole life. They have been influential and supportive at every step, including keeping my sights set high. I have them especially to thank for making me apply to Yale and then enabling me to attend. Their value of my education makes completion of this degree all the sweeter.

Perhaps the strongest direct source of support I received as a student was from my friend and officemate Pablo. A brilliant scientist, a sympathetic poet, a loyal friend, no other classmate has been so giving and supportive in my frequent times of darkness. I most certainly intend to make good on my promise to visit him and his family in Sevilla.

And lastly, none of this would be possible without my literal life blood, my immediate family. Clearly my largest fans, my Mom, Dad, and sister Amy have been present for the best and worst of me. Through tears of depression and elation at all hours of the day and night, they have counseled and celebrated this degree's entirety. I thank them for their constant love and support.

# Contents

<b>1</b>	<b>Introduction &amp; Overview</b>	<b>11</b>
<b>2</b>	<b>Prediction and Predictability: Background</b>	<b>19</b>
2.1	Weather prediction . . . . .	20
2.1.1	The approach — NWP . . . . .	20
2.1.2	The components of NWP . . . . .	22
2.2	Nonlinear dynamics — a geometric perspective . . . . .	24
2.2.1	A review of the basics . . . . .	24
2.2.2	Dissipative systems and attractors . . . . .	28
2.2.3	Chaotic systems and strange attractors . . . . .	30
2.2.4	Prediction: geometric interpretation . . . . .	33
2.3	Nonlinear dynamics & predictability . . . . .	34
2.3.1	Prediction: the correct probabilistic approach . . . . .	34
2.3.2	Approximate probabilistic approaches to prediction . . . . .	40
2.3.3	Practical probabilistic approaches to prediction . . . . .	43
2.4	Nonlinear dynamical approach to NWP . . . . .	58
2.4.1	The relevance of attractors . . . . .	58
2.4.2	Operational approaches . . . . .	59
2.4.3	“Physics-based” approaches . . . . .	62
2.5	Background summary . . . . .	64
<b>3</b>	<b>The Impacts of Nonlinearity and Non-Gaussianity in Ensemble-Based Data Assimilation</b>	<b>67</b>
3.1	Introduction . . . . .	67
3.2	Filter formulations . . . . .	69
3.3	Geometric interpretation of ensemble filter updating . . . . .	72
3.3.1	Update kinematics: one-dimensional example . . . . .	73

3.3.2	Update dynamics: the Ikeda system . . . . .	77
3.4	Ensemble filter behavior in higher dimensioned systems . . . . .	89
3.4.1	Lorenz latitude circle model . . . . .	90
3.4.2	Double-gyre barotropic ocean model . . . . .	91
3.5	Conclusions . . . . .	93
<b>4</b>	<b>Position Errors &amp; Lagrangian Estimation</b>	<b>97</b>
4.1	Features and position errors . . . . .	98
4.1.1	Position errors . . . . .	99
4.1.2	Historical treatment of position errors . . . . .	101
4.1.3	Contour analysis . . . . .	104
4.1.4	Feature calibration and alignment . . . . .	107
4.1.5	Tropical cyclone prediction . . . . .	107
4.1.6	Position error conclusion . . . . .	110
4.2	The point vortex model & Lagrangian estimation . . . . .	111
4.2.1	The point vortex model . . . . .	111
4.2.2	Previous point vortex predictability studies . . . . .	114
4.2.3	Ensemble-based Kalman filtering . . . . .	116
4.2.4	Comparing Eulerian and Lagrangian state vectors . . . . .	120
4.2.5	Point vortex summary . . . . .	139
4.3	Chapter summary . . . . .	139
<b>5</b>	<b>Two-Step Ensemble-Based Data Assimilation</b>	<b>141</b>
5.1	Introduction . . . . .	142
5.2	Error models . . . . .	145
5.2.1	Conceptual error representations . . . . .	146
5.2.2	Non-uniqueness of error models . . . . .	148
5.2.3	A mixed alignment and additive error model . . . . .	150
5.3	Estimation and error models . . . . .	154
5.3.1	Kalman filtering: basic assumptions and machinery . . . . .	155
5.3.2	Ensemble-based filtering . . . . .	156
5.3.3	The Kalman filter as a repositioning device . . . . .	158
5.3.4	Examples of filter updates . . . . .	162
5.4	A two-step filter approach . . . . .	174
5.5	Observation system simulation experiments . . . . .	178
5.5.1	Experimental set-up . . . . .	178

5.5.2	The failure of the tangent linear hypothesis . . . . .	179
5.5.3	Ensemble updates . . . . .	180
5.6	Conclusion . . . . .	185
<b>6</b>	<b>Extensions to multiple features and multiple spatial dimensions</b>	<b>189</b>
6.1	An augmented state vector approach . . . . .	191
6.1.1	Augmenting mechanics and statistics . . . . .	192
6.1.2	Augmented state vectors and state estimation . . . . .	193
6.2	Basis estimation and alignment schemes . . . . .	196
6.2.1	Basis estimation . . . . .	197
6.2.2	Alignment schemes . . . . .	200
6.2.3	Feature identification & tracking . . . . .	210
6.2.4	Feature correspondence across the ensemble . . . . .	211
6.3	OSSEs . . . . .	211
6.3.1	A barotropic model . . . . .	211
6.3.2	The point vortex model revisited . . . . .	213
6.3.3	A barotropically unstable jet . . . . .	231
6.4	Summary . . . . .	249
<b>7</b>	<b>Conclusions &amp; Future Work</b>	<b>251</b>
7.1	Chapter summaries . . . . .	251
7.2	Overall summary . . . . .	254
7.3	Continuing directions . . . . .	255
7.3.1	Logical extensions . . . . .	256
7.3.2	Possible applications . . . . .	256
7.3.3	Ensemble compositing . . . . .	257
7.3.4	Reusing observations . . . . .	258
<b>A</b>	<b>Skewness and kurtosis</b>	<b>261</b>
A.1	Definitions . . . . .	261
A.2	Maximum Values . . . . .	262
<b>B</b>	<b>Expected total error statistics</b>	<b>263</b>
<b>C</b>	<b>Lorenz latitude circle model</b>	<b>265</b>
<b>D</b>	<b>Forms of numerical dissipation</b>	<b>267</b>

D.1 Wavenumber filter . . . . .	267
D.2 Biharmonic hyperdiffusion . . . . .	268
<b>E Doubly periodic point vortex model</b>	<b>269</b>

# Chapter 1

## Introduction & Overview

This thesis focuses on state estimation, or data assimilation as it is often called. In particular, it focuses on implementable methods of state estimation suitable for use in numerical weather prediction applications. As shall be described later, there is, in fact, a theoretically correct approach to probabilistic state estimation in regimes of nonlinear error growth, but it is unavailable for use in all but the simplest of problems. Due to fairly rigid pragmatic constraints, essentially all existing methods of implementable state estimation, including the state-of-the-art methods, are forced to assume that errors remain in regimes of linear error growth.<sup>1</sup> However, a large part of what makes state estimation an interesting pursuit is its use in systems with nonlinear, chaotic dynamics, and such dynamics can easily allow regimes of nonlinear error growth, which, in turn, can produce distributions of non-Gaussian uncertainty.<sup>2</sup> The ability of systems of interest to violate the assumptions of available methods of state estimation poses a potential problem, and we wish to better understand this problem, particularly its impacts and its relevance to numerical weather prediction.

Very broadly, in the thesis we aim to push the state-of-the-art methods of state estimation past their limits of validity and see how and when they fail. Then, having learned this, we wish to see if these limits can be extended a bit further based on physical intuition gained from an important class of meteorological problems, namely the evolution of coherent features. In particular, we seek the answers to the following questions:

- 1) How do the state-of-the-art methods of state estimation fail in the face of nonlinear error growth and non-Gaussian uncertainty?

---

<sup>1</sup>To be clear about this, a regime of linear error growth is one where the error dynamics are approximately linear, meaning that the error amplitudes grow approximately exponentially.

<sup>2</sup>In general, some nonlinearity is required to produce non-Gaussian uncertainty; however, the nonlinearity need not be restricted to the system dynamics. For example, a perfectly linear system forced by multiplicative noise can also produce non-Gaussian uncertainty.

- 2) How can we successfully perform state estimation when mis-positioned coherent features are present?

To answer the first question, we consider quite generic sources of non-Gaussian uncertainty through a hierarchy of model complexity. We then argue that mis-positioned coherent features constitute a specific, geophysically relevant source of non-Gaussian uncertainty.

We find we are able to extend the limits of implementable state estimation methods in some scenarios where mis-positioned coherent features are present due to an available redefinition based on an understanding of the phenomenology. The redefinition is primarily based on the differences between two equivalent fluid descriptions, Eulerian (i.e., discrete gridpoint-based) and Lagrangian (i.e., fluid parcel-based) descriptions. Specifically, distributions of Gaussian uncertainty in one description can be non-Gaussian in the corresponding, alternate description, and knowledge of this suggests an approach to extending the limits of validity for implementable methods of state estimation. As we shall see, Gaussian feature position errors (Lagrangian description) can easily have corresponding non-Gaussian amplitude errors (Eulerian description). Coherent features are nearly ubiquitous in geophysical fluids, and their state estimation and prediction represent a frontier in predictability science, particularly at the mesoscale.

Due to the practical constraints limiting the scope of implementable methods of state estimation, an alternative view of our above questioning is: how can we maximize the utility of the state-of-the-art methods of state estimation? Accordingly, we aim to identify the extent to which existing state estimation frameworks can be utilized and augmented to be successful under the influence of the specific and relevant source of non-Gaussian uncertainty owing to mis-positioned coherent features. And though the motivating aim is implementation in the operational numerical weather prediction setting, in this thesis we have taken a hierarchical approach and have chosen to learn from manageable, understandable examples.

Within the thesis we rely almost exclusively on ensemble-based methods, or Monte Carlo methods. We do this for several reasons. The first is that we have interest in performing probabilistic state estimation and predictions, and if one can maintain an ensemble of states that are statistically indistinguishable from truth, then ensemble-based methods constitute the closest available approximation to the correct probabilistic approach. And because of this, ensemble-based methods afford estimates of flow-dependent uncertainty, which is an important and information-rich advance beyond currently operational methods of state estimation. Also, ensemble-based methods constitute the state-of-the-art for implementable methods of state estimation, and we therefore focus on them in our treatment of the above questions.

While our immediate interests lie in improving state estimation for numerical weather prediction purposes, our over-arching goal is shared by all dynamical meteorologists, and that is to ultimately understand the atmosphere better. State estimation is merely a component in the grand, interconnected scheme of hypothesis testing via the scientific method, and, as such, should be viewed as a tool in the same way that model development, observational development, verification methods, ensemble assessment methods, and similar fields should be as well. When all the components are assembled and synchronized, we have the ability to test our furthest understanding of the atmosphere through making and verifying predictions of the system. State estimation is the interface by which the atmosphere, via observations, can ultimately affect our models of it.

### **Numerical weather prediction's role in atmospheric science**

Understanding the evolution of the earth's atmosphere is the broad goal of dynamical meteorology. The reach of this understanding has grown in accordance with the development of fluid dynamics as a whole over the past several centuries. After the physics thought to govern a fluid's evolution were learned and written down, the bulk of which had occurred by the mid-nineteenth century, the quest for understanding had merely begun. The physics eludes analytical, closed-form solutions, and hence remains opaque to traditional methods of analysis. At best, only academic, limiting examples can be treated, and non-dimensionalization and scaling analysis performed. To be sure, our understanding of the atmosphere's mean state and its instabilities has increased (though is still far from solved) and their characterizations have improved; however, we still lack understanding rich enough to predict the details of transient events, that is, weather. Fortunately, knowledge of the governing equations themselves provides a means to approximating their solution.

The year 2004 marks the one hundredth anniversary of an important idea put forth by Vilhelm Bjerknes, one of the founding fathers of modern dynamical meteorology. In his classic 1904 paper entitled, "The problem of weather forecasting as a problem in mechanics and physics," Bjerknes submitted the claim that the atmosphere is merely a complicated system whose evolution is, in principle, able to be "precalculated" based on knowledge of the governing physics and an initial condition. Within this paper, Bjerknes was apparently the first to offer the laws of hydrodynamics and thermodynamics as the governing equations for the atmosphere. By treating weather prediction as an initial value problem, Bjerknes linked meteorology to the broader realms of mathematical physics and nonlinear dynamics. Bjerknes later referred to weather prediction as "the ultimate problem in meteorology" (1911), thereby opening the door to a science of *predictability*.

While Bjerknes evidently understood the enormous complexity of the problem logistically, there is much that has been brought to bear on the matter of predicting complicated systems in the century since his original paper, due in large part to the introduction of the digital computer. Indeed, the dawn of numerical weather prediction has brought Bjerknes's vision closer to fruition; yet simultaneously, computers have given scientists the means to appreciate and assess predictability in general in nonlinear systems. What they find is that, conceptually, weather prediction is as straightforward as Bjerknes states it to be, but there are a great many issues of practicality and implementability (beyond logistics) of which he could not have been aware.

This issue of practicality is not to be taken lightly. As we shall discuss, the treatment of the atmosphere as a nonlinear dynamical system implies a probabilistic approach to its state estimation and prediction; a theoretically correct approach to probabilistic state estimation and prediction exists, so in some respects one might dismiss weather prediction as a "solved problem." However, once one begins to examine the correct method and understand its issues, one immediately finds that there are limits to its usefulness set by *a*) our computational maturity, and *b*) the raw amount of information required to close the prediction problem. These remain difficult problems of a practical nature whose existence impedes delivery of the very product upon which many millions rely. Of course, weather forecasts are issued, and will continue to be issued, regardless of their scientific content and rigor, but there is still much that science can offer toward improving these forecasts, where "improve" both refers to the variously assessed skill scores that are used to judge a prediction, and more generally in the sense of offering sound procedural methods.

### **The importance of probabilism**

Though many, particularly non-scientists, take a *deterministic* point of view to weather prediction, this thesis assumes a wholly *probabilistic* view. It has become clear that weather prediction will always suffer from unavoidable errors. These errors will prevent us from ever knowing the exact state of the atmosphere (i.e., the true state). In fact, all specified atmospheric states are merely estimates of the true state, and therefore, to be meaningful, should be accompanied by measures of their uncertainty. In simple laboratory measurements, the standard error (i.e., the error bar) is usually sufficient, but in the realm of nonlinear dynamical systems, in which Bjerknes envisioned weather prediction belonging, the notion of uncertainty is best kept in its most general form, namely the probability density function (PDF, also probability distribution function, the difference being that densities are differentiable, whereas distributions need not be).

The aforementioned theoretically correct method to probabilistic state estimation and prediction has a clear meaning in state space — a system’s attractor is its climatological PDF, and available observations serve to sharpen our initial condition estimate as a conditional PDF. However, since the correct method cannot be implemented, approximations must be made. To date, most implemented methods treat only a single state and possibly its ensemble-based, flow-dependent covariance. These can be interpreted as the first two central moments of a PDF, and given certain simplifying assumptions these two moments can completely characterize the PDF. However, the nonlinearity of atmospheric dynamics readily allows PDFs to violate these assumptions, and thus implementable methods can fail.

The correct method is based geometrically on the system’s attractor in its state space. An attractor encompasses all that is known about the system dynamics, it is the graph of the system’s solution set. Once one makes the assumption that the state’s PDF is Gaussian, one potentially discards the rich, non-Gaussian (fractal) structure of the attractor. This in turn can result in the loss of much of the dynamical information available that is potentially important for state estimation in the nonlinear regime. This lost information is mainly a consequence of implementation — the theoretically correct approach does not discard this information, and is consequently too unwieldy to be useful. Therefore, to recover this lost information important for the estimation process, one might turn from *state space* considerations (where a state estimate is a point) to *model space* (or physical space) considerations (where a state estimate is a discretized fluid field), with the intent of introducing dynamical information about the current state into implementable methods of state estimation. There are many approaches to this that one might try; one key idea explored in this thesis concerns redefinitions of an estimate’s errors based on some understanding of the estimate’s current state. As we shall see, while errors in state space have a clear meaning (given a suitably defined norm), errors in model space are non-unique. And as estimation methods can be seen as error reduction methods, estimation methods applied in model space are also non-unique. Estimation in model space requires specification of an *error model*, a statement of how states are expected to differ from truth. Most implementable state estimation schemes assume that errors are additive, and we ultimately exploit the notion of alternative error models as a means to naturally transform the estimation problem into a two-step approach when considering the non-Gaussianity imparted by the mis-positioning of tight, coherent features within a fluid domain.

Depending on one’s notion of what constitutes a feature, coherent features are arguably ubiquitous within the atmospheric state. An archetypal atmospheric feature could be the

tropical cyclone, but really any localized concentration of a meteorologically relevant quantity could constitute a feature. Resolved features in numerical weather prediction pose the possibility of being mis-positioned, and this, we show, can lead to significant non-Gaussian PDFs, thereby violating the implementable methods' assumptions. Correctly forecasting features is important, but it can derail traditional methods of state estimation. Strictly, to predict features correctly, one needs a state estimation scheme capable of working with non-Gaussian PDFs. Given that none seem likely implementable, features will remain able to undermine traditional, implementable methods. *Hence, we are looking ahead to practical methods of state estimation able to correct the expected non-Gaussianity due to the presence of features.*

## Outline

We seek to address the issues raised above and to ultimately offer a plausibly implementable approach to state estimation capable of treating the non-Gaussian error PDFs that arise in specific, physically relevant situations. To this end, we have compiled this thesis. It is arranged as follows:

Chapter 2 provides background information on some integral concepts in predictability and nonlinear dynamics that the uninformed reader will likely require to appreciate the material to follow in subsequent chapters. There is a general emphasis on the geometric interpretation of system dynamics, so that the theoretically correct approach to probabilistic state estimation and prediction can be demystified. We then review the correct approach's relationship with the large majority of implementable methods. It is quite possible the informed reader can skim this chapter or skip it entirely.

Chapter 3 examines the impacts of nonlinear error growth and non-Gaussian uncertainty on the state-of-the-art methods of implementable state estimation. This chapter considers rather generic sources of non-Gaussian uncertainty, and it contains the bulk of the work toward addressing our first posed question above of how and when these methods fail. We develop a geometric understanding of the methods' updates in visualizable state spaces, and then go to lengths to show our understanding is consistent with the methods' behavior through a hierarchy of model complexity.

Whereas chapter 3 considers generic sources of non-Gaussian uncertainty, chapter 4 hones its focus to the specific, geophysically relevant source of non-Gaussian uncertainty owing to mis-positioned coherent features. After arguing that position errors are indeed relevant to numerical weather prediction, we then explore methods of state estimation encased in a Lagrangian (i.e., fluid parcel-based) framework, which we expect will ultimately

not be very useful for numerical weather prediction applications.

In chapter 5 we shift the focus back to methods of state estimation encased in an Eulerian (i.e., gridpoint-based) framework, which we expect will prove relevant to numerical weather prediction. Within this thesis we develop the groundwork for a two-step approach to state estimation suitable for use with coherent features that relies only implementable methods of state estimation. This chapter outlines the basics of this proposed approach as encapsulated by alternative error model assumptions.

Chapter 6 extends the work presented in chapter 5 to more realistic settings. In particular, we focus on extensions to systems exhibiting multiple simultaneous features and which operate in multiple spatial dimensions. We present several proof-of-concept experiments to establish plausibility and to demonstrate the expected successful extension into regimes of nonlinear error growth. The necessity and importance of an alignment scheme is developed and stressed.

Chapter 7 then concludes with a thesis summary and some overall conclusions. We also suggest some possible avenues of further research on the matter.



## Chapter 2

# Prediction and Predictability: Background

This chapter serves first and foremost to orient readers. There are no new results *per se*, though we have generated a few examples to help communicate points. The over-riding goal is to reinterpret prediction and state estimation informed by knowledge gained from studies of nonlinear dynamical systems. The most important points we want to communicate are: a) prediction of chaotic systems ought to be *probabilistic*, b) nearly all extant prediction systems suitable for spatially extended chaotic systems assume that error growth follows linear dynamics and retain a normal distribution, and c) “state-of-the-art” weather prediction systems are also based on these rather restrictive assumptions. The informed reader may already appreciate these points, but we make no such assumptions, particularly as predictability is not (yet) a core concept in the education of atmospheric scientists. As appreciating these points requires a basic understanding of nonlinear dynamics, we take time to review these basics. Our review is by no means exhaustive, it is aimed at fostering a geometric perspective on dynamical systems with emphasis placed on state spaces and attractors. We make an effort to note where each reviewed concept is encountered in the rest of the thesis. It is the intent of the review to inform readers sufficiently so that we may “demystify” the processes of state estimation and prediction, and so that we may give some insight into why they can be difficult ventures.

As this thesis mainly focuses on state estimation, we also take time to review its basics. We begin with the theoretically correct approach, namely Bayes rule, and discuss the limitations of its implementation in all but the simplest problems. We then discuss various simplifications of the correct approach, ultimately arriving at the practical, implementable state estimation methods used currently. We wish to impress that the two most limiting

factors facing state estimation schemes, namely computation and information, are difficult to surmount, and that, consequently, simplifying assumptions of linear error growth and normal error distributions are pervasive. We take time to derive and discuss how Bayes rule reduces to the Kalman filter, or equivalently, least squares, since approximations and extensions to the Kalman filter are the focus of most of the work in this thesis. Also, we establish a notation and nomenclature to which we adhere for the rest of the thesis.

We also briefly discuss the current state estimation schemes being used in major operation forecasting centers around the world. It is our intent that our work could find application within the realm of numerical weather prediction (NWP), but this thesis is restricted to relatively simple models and examples. Working with simple systems allows an obvious computational advantage for our experiments, but even had we a dedicated supercomputer facility for our experiments, we would still begin by examining these simple systems as they allow greater clarity into the scientific issues. Hence, we often approach problems from a hierarchical perspective in an attempt to demonstrate that insight gained from simple models is consistent with what is observed in larger models. The philosophy of this approach is stated quite well by Lenny Smith: “Although it is unreasonable to expect solutions to low-dimensional problems to generalize to million dimensional spaces, so too it is unlikely that problems identified in the simplified models will vanish in operational models” (p. 2487, Smith 2002).

The informed reader can safely skip this chapter. For the less informed reader, we have broken the material into short, entitled subsections, and have emphasized key sentences and concepts to facilitate browsing and skimming as necessary.

## 2.1 Weather prediction

### 2.1.1 The approach — NWP

As envisioned by Bjerknes, a sensible approach to predicting the weather would be to integrate the appropriate governing physics. Most of the physics governing the dynamical evolution of the earth’s atmospheric state was known to humans by the mid-nineteenth century. Though the atmosphere is a primarily radiatively forced system, its internal fluid dynamics are generalizations of classical mechanics to fluid continua. Depending on the level of detail considered (e.g., active atmospheric chemistry), most operational weather forecasting approaches consider seven distinct variables in time over a spatial domain: the Navier-Stokes equations govern the time evolution of the components of momentum,  $u$ ,  $v$ ,

and  $w$  (though sometimes the hydrostatic approximation is made in the vertical, leaving  $w$  to be diagnosed); the first law of thermodynamics governs the time evolution of temperature,  $T$ ; the equation of mass conservation governs the time evolution of density,  $\rho$ ; a second conservation equation governs the time evolution of water substance,  $q$ ; and an equation of state allows diagnosis of the pressure,  $p$ . There are many approximations made in formulating these equations, yet scientists believe these approximations are minor enough that the equations are fairly accurate in modeling the atmospheric state's actual evolution. That these equations are known is somewhat remarkable since weather prediction is subsumed in their solution. Indeed, an analytical, closed-form solution to this system of nonlinear, coupled partial differential equations (PDEs) would yield an attractive means to achieve weather prediction. However, like most nonlinear PDEs, a general analytical solution remains elusive. In fact, one of the Clay Mathematics Institute's Millenium Prizes offers an enticing one million dollars to any scientist who can (merely) demonstrate properties of solutions' smoothness to the Navier-Stokes equations alone.

The possibility of analytical solutions aside, knowledge of the governing equations allows for other, more approximate attempts at their solution. The most attractive by far has been the application of numerical methods. In numerical weather prediction (NWP), the continuous domain is discretized in space and time in some fashion and an initial condition is stepped forward. There are two broad components controlling the accuracy of NWP: the accuracy of the model and its discretization, and the accuracy of the initial condition. These components are intimately related though we do our best to isolate them: this thesis focuses on the accuracy of initial conditions and largely assumes that the models used are perfectly accurate (including all parameters and boundary conditions). However, before moving forward, we briefly address issues associated with using inaccurate models.

NWP's required discretization has the consequence of transforming the system of *differential* equations, intractable to advanced solution methods based on calculus, into a system of *difference* equations whose solution can often be managed by algebra. For consistent discretization schemes, the resulting algebraic forms can be shown to approach their differential forms by allowing the spacing between discrete elements, the model's "resolution," to approach zero and then evaluating the appropriate limits. The introduction of a resolution necessarily introduces further approximations to the solution of the equations, both in that discretizations are simply not equal to differentials, and also by necessitating parameterization of processes that occur on scales smaller than the resolved scales. Hence, as is perhaps obvious, even if the original differential equations are perfect in the sense that they exactly govern the true evolution of the earth's atmosphere (which we know they do not),

solving them numerically instantly introduces imperfections (i.e., errors) to the prediction solutions.

Discretization of the governing equations allows for a tractable approach to weather prediction; however, a consequence is that one must explicitly represent each variable at each discretized location in space by a number of some given accuracy. When one considers the full equations (i.e., 7 variables) represented over the globe with fine resolution in the horizontal and vertical, the number of numbers required to represent the state becomes huge. State-of-the-art atmospheric state estimates have sizes somewhere between  $10^7$  and  $10^8$ , a computationally formidable entity to treat. Because the discretization spacing is inversely proportional to the discretized equations' accuracy, there is an obvious trade-off between practical implementation and obtaining realistic predictions: as the resolution approaches a continuous representation of the fluid, the problem becomes computationally intractable. Hence, in order to produce forecasts conditioned on the governing equations, one must live with errors introduced by numerical methods on top of any existing errors owing to approximations made in originally formulating the governing equations. These errors are broadly grouped together as sources of "model error." Model error is unavoidable when predicting phenomena in nature. Model error is a difficult problem and an active area of research. Model error can lead to initial condition error after one timestep. In this thesis we focus on so-called "perfect models," where initial conditions without error will yield predictions without error. We do this with the belief that if we cannot understand the ideal scenario, then we have no hope understanding more complicated, realistic cases.

### 2.1.2 The components of NWP

Numerical weather prediction as it exists today is comprised of three major components: a forecast model, observations, and a data assimilation system. The forecast model evolves a specified state forward in time. At best, the forecast model would be the true solution to the continuous equations, but really any approximation thereof may serve as the forecast model. *The important thing to realize is that all models are acceptable, and all are only approximations of truth. More valuable models have smaller errors attributable to their given approximations.*

Observations are measurements of truth, and may therefore be termed instrumental estimates. Observations do not rely upon any approximations made in a model formulation, but they pose the problem of being an incomplete estimate, that is, *the atmospheric state is only ever partially observed.* Scientists observe the atmospheric state directly and indirectly through a variety of observational platforms. Direct observations are measurements made

for one of the variables predicted by the forecast model, and indirect observations are measurements made of quantities related, possibly nonlinearly, to the model's prognostic variables. For example, a thermometer directly measures temperature, whereas a satellite measures radiance, an integrated measure of temperature. As explained by Morss (1998), there are four types of observational platforms: surface observation stations and radiosonde release sites worldwide, such as those run by the National Weather Service, help constitute *fixed platform* observations; commercial aircraft and ocean liners are examples of *platforms of opportunity*; satellites are *remote sensing platforms*; and as explored by Morss and others, a fourth type of platform receiving more attention is the possibility of *targeted observations*, that is, observations made as dictated by the "dynamics of the day." Observations from all platforms provide estimates of the true atmospheric state at given locations and at given times. All observations have errors, and observational instruments should, in principle, come with estimates of their expected uncertainties. It should be noted that an important issue in operational NWP is that of observational quality control (e.g., Dee et al. 2001). While this could constitute a component of NWP on its own right, we include it within the process of observational retrieval, if only to isolate the third component, the main focus of this thesis.

The third component of NWP is data assimilation (DA). DA is the main focus of this thesis. "Data assimilation" is a term popular in meteorology, but it is just another name for state estimation or signal processing. There are generally two broad solution methods within state estimation, each of which can be applied to the two broad types of state estimation problems. The solution methods are so-called sequential methods and variational methods, the former constituting a direct solution and the latter constituting an iterative solution. The two types of state estimation problems are filtering and smoothing. Filtering seeks the best state estimate at a certain time using all available information up to and including that time, whereas smoothing seeks the best estimate at a certain time or over a certain time interval using all available information total, even possibly observations taken after the time of interest. Smoothing is useful in problems where one desires a continuous record of consistent state estimates for postmortem analysis and diagnosis (e.g., for energy budgets); however, the end product of DA in an NWP context is the initial state used to numerically solve the equations, and therefore filtering is the more relevant state estimation problem to consider here. Unless otherwise noted, reference to data assimilation within this thesis will specifically mean filtering.

The initial state constitutes the best estimate of the true atmospheric state for a given time. DA is most often posed as a Bayesian estimation procedure, meaning that a prior

estimate is considered. The prior estimate is typically a short term forecast (e.g., a 6 hour forecast) from the last time DA was performed. Fundamentally, DA serves to combine two information sources, the prior estimate and newly available observations, to achieve its best estimate. However, because there are typically many fewer observations available than the dimension of the atmospheric state (i.e., the  $10^7$  referred to above), another common interpretation is that DA serves to correct the prior estimate based on new information provided by the observations that have become available since that short term forecast was launched. Either way, the combining process or correcting process can be carried out in a variety of ways, but most serve to weight the two information sources by their respective uncertainties. The specifics of this process as well as a discussion of theoretically “correct” procedures to accomplish this will be presented shortly. The importance of DA will become apparent after an introduction to studies of nonlinear dynamical systems and a discussion of some of the pertinent lessons learned therein.

## 2.2 Nonlinear dynamics — a geometric perspective

Since Bjerknes’s original idea, it was the advent of computers that eventually allowed for realistic NWP. Computers have also allowed for studies of predictability. Before introducing the important results learned, we first take time to review the basics of dynamical systems so that we can reinterpret NWP as Bjerknes (1904) conceived it, only armed with a century’s more knowledge. In this section we briefly review the basics of nonlinear dynamics with a focus on the geometry of state space and attractors. We introduce the notion of chaos and strange attractors. The informed reader should be able to skip to section 2.2.4 on the geometric interpretation of prediction.

### 2.2.1 A review of the basics

In this subsection we briefly review the important concepts of dynamical systems. In particular, we discuss the notion a system’s state space.

#### State spaces

A *system* is defined by specified governing equations of motion, as long as they admit solutions (e.g., are not under-determined). By convention, all model parameters are part of the system. A system can be comprised of ordinary differential equations (ODEs), partial differential equations (PDEs), or difference equations; it is the presence of a time derivative

or time step that qualifies the system as a dynamical system. A given solution to the system defines a *state*. By convention, a state is stored as a column vector,  $\mathbf{x}$ , of dimension  $n$ . Although most PDEs, like the governing equations for the atmosphere, have continuous functions as their solution, it is still convention that a state is a column vector, however, it should be kept in mind that a column vector representation implies that any continuous functions have already been discretized in some fashion. If analytical solutions were known, then one could in principle treat them as states.

A system's *state space* is the  $n$ -dimensional space where the system's solutions exist. A state is thus a point in state space with one state element per orthogonal axis.<sup>1</sup> Distance within a state space is only defined in reference to an appropriate norm (e.g., the  $\ell^2$  norm). Much of the work and many of the breakthroughs in nonlinear dynamics have come from *geometric* considerations of and within a system's state space. Clearly, in order to be able to visualize a system's state space, thereby taking full advantage of its geometric interpretation, it must be of low dimension. Luckily, there are many illuminating low dimensional examples that have been considered in the literature. *Whether the insights gained from studying low dimensional systems extend directly to the hugely dimensioned systems of geophysical interest is an open question; however, the insights gained do provide a framework in which to think, and they in some sense demystify the mechanics of prediction and data assimilation.* In this thesis, when trying to extend lessons learned from low dimensional systems to geophysically relevant systems, we employ a hierarchical approach to try to at least establish consistency through the tiers of complexity.

The real utility of state space considerations comes forth in analysis of a state's evolution. The evolution of a state under a given system dynamics is a trajectory through state space. It is assumed throughout this thesis and much of the nonlinear dynamics literature that the system dynamics are *deterministic*. This disallows evolving states' trajectories to intersect in state space. By specifying  $\mathbf{x}$ , one precisely places the state in its state space. The system dynamics are assumed to adhere to a causality principle, that is, future and past evolution are determined uniquely by knowledge of the present state only. If the dynamics have "memory," such that knowledge of a past state will influence its evolution, or if the dynamics are nonautonomous (i.e., depends on time explicitly), then the system and state can always be generalized to include this extra dependence, thereby rendering the dynamics Markov again. There is a growing literature that considers *stochastic* dynamics. The notion of a state space is less useful in stochastic systems, but prediction and data assimilation

---

<sup>1</sup>Without discretization, a system of PDEs defines a "function space," the space in which all continuous solutions over a specified interval exist, thus making a function a point in function space.

still proceed by the same mechanics.

The state spaces of simple one and two-dimensional systems will be considered heavily in section 3.3. The other models considered within the thesis have state spaces too large to visualize directly, but state space concerns always underlie their state estimation.

### **Flows versus maps**

One broad classification of dynamical systems is that of flows versus maps. Flows arise from differential equations and have generally continuous trajectories through their state spaces. Maps spring from difference equations, and their iterative nature traces system evolution through discontinuous, discrete steps. In practice, flows can easily be rendered as maps. The solution method of most nonlinear differential equations is to attack them numerically as discussed in section 2.1.1, and this transforms differential equations into difference equations (i.e., maps). Also, the first glimpse of chaotic dynamical behavior was caught by Poincaré when he made his famous Poincaré section in his studies of the three-body problem. He considered a plane in a higher dimensional state space and collected the points where evolving continuous trajectories intersected his plane of interest. The resulting Poincaré section is thus a discontinuous mapping from each successive crossing of the plane, even though the full system is a flow. Poincaré sections are two-dimensional mappings in that one considers a plane in state space. Lorenz (1963) devised a one-dimensional mapping by plotting a flow's relative maximum value versus its immediately previous relative maximum value. In this so-called first return map, or sometimes just Lorenz map, one can sometimes succinctly collapse seemingly erratic evolutionary behavior onto a nearly ordered structure. Poincaré sections and first return maps have proven a very useful approach to considering evolution in state spaces with dimension too large to fully visualize.

The Ikeda system studied in section 3.3 is an example of a map. Most other systems considered in this thesis are conceptually flows, however, their numerical solution method requires discretization and thus renders them all effective mappings.

### **Conservative versus dissipative**

Another broad classification of dynamical systems is that of conservative versus dissipative systems. Conservative systems are frictionless, and are thus in many ways idealized. Physical systems, by and large, are dissipative. A general approach to describing frictionless dynamical systems is through Hamiltonian mechanics. Without presenting too much detail, a frictionless system allows formulation of a Hamiltonian function,  $H$ , which is similar to the system's total energy (and under some circumstances is identical to the total energy).

If the Hamiltonian function is autonomous, then the system is said to be conservative, and trajectories through its state space follow isopleths of constant  $H$ . Being conservative means that the system dynamics are “reversible,” that is, the system behaves identically under the transformation  $t' = -t$ . Also, evolution through state space conserves area (or more generally, “hypervolume”), meaning that neighboring trajectories do not depart from one another too quickly, or in the case of dilating trajectories in one direction, it is certain that there are contracting trajectories in an orthogonal, compensating direction. In contrast, dissipative systems are irreversible and have contracting areas in their state space over evolution.

In this thesis we consider the state estimation and prediction of both conservative and dissipative systems. Though only the latter have attractors (to be defined below), both types of systems have state spaces, and both types of systems can exhibit the sensitivity to initial conditions that ultimately make filtering difficult. The point vortex system considered in sections 4.2 and 6.3.2 and the Korteweg-de Vries equation considered throughout chapter 5 are examples of conservative systems, whereas the Ikeda map, the Lorenz latitude circle model, and the double-gyre barotropic ocean model considered in the next chapter are examples of dissipative systems. The doubly periodic barotropic model used in chapter 6 includes numerical dissipation (so-called hyperviscosity), and so is a dissipative model even though we use it to mimic a conservative point vortex model.

## Stability

One of the main interests in studies of nonlinear dynamical systems is to characterize a system’s stability. This is a very large topic — we cannot and should not review everything. We merely mention that the geometric treatment of state space gives insight into the existence and behavior of a system’s stability. A system’s state space has special points, equilibrium points called *fixed points*, where the dynamics predict no motion through state space. Stable means that perturbations from the fixed point will meet a restoring force that tends to return them to the fixed point, and unstable means that perturbations from the fixed point will meet a force that tends to draw them farther from the fixed point. Linear stability examines the stability of small perturbations, whereas nonlinear stability examines the stability of conditional scenarios and finite-amplitude perturbations.

We note that a system’s stability is in general state dependent. Stated another way, the sensitivity of each type of fixed point to perturbations is in general different. As errors can be considered perturbations in a system’s state space, this notion of state dependent stability foreshadows the concept of state dependent error growth (to be discussed in the next

subsection). This issue is central to our choice of using ensemble-based data assimilation throughout the thesis (to be discussed further in section 2.3.3).

## 2.2.2 Dissipative systems and attractors

As most physical systems are dissipative, it is useful to contrast the qualitative differences between a conservative system’s state space and that of a dissipative system. The main difference is the emergence of an attractor, to be defined below.

### Attractors

One of the most important aspects of dissipative systems is the existence of attracting sets in state space, also simply called *attractors*. States already on the attractor stay on the attractor, and states off the attractor (but still within its “basin of attraction”) evolve toward the attractor. This is a powerful constraint, and it makes the geometric treatment of dissipative dynamical systems a fruitful pursuit.

Strictly, an attractor is a point set (sometimes referred to as a “manifold,” though an attractor need not be smooth) in a system’s state space that defines all the states the system could ever freely occupy (as  $t \rightarrow \infty$ ). Due to the contraction of state space areas for dissipative systems, the attractor has zero (hyper)volume in the  $n$ -dimensional state space. As noted by Strogatz (1994), precisely defining an attractor is difficult, though he does offer a sensible working definition: an attractor is a closed set in a system’s state space that a.) is invariant, b.) attracts an open set of initial conditions, and c.) is minimal. The first condition means that state space trajectories within an attractor stay in the attractor; the second condition means that trajectories begun outside the attractor (but still within the basin of attraction) tend toward the attractor; and the third condition means that there is no subset of the attractor that satisfies the first two conditions on its own. A given system can readily have multiple simultaneous attractors.

### Lyapunov exponents

*Lyapunov exponents* characterize the attractor-wide average linear stability — if a system has a positive Lyapunov exponent, then in the limit of infinite time, an arbitrarily small initial perturbation added to *any* point on a system’s attractor will come to grow at a rate given by the largest Lyapunov exponent, and the perturbation will grow into the structure determined by that exponent’s corresponding Lyapunov vector. Lyapunov exponents and vectors can be thought of as the normal modes of an attractor, that is, the eigenvalues and

eigenvectors of the attractor-wide averaged Jacobian matrix (linearized dynamics). Because of the attractor-wide averaging (or alternatively, the infinite time allowed for evolution), Lyapunov exponents and vectors are state independent. *Because of this state independence, Lyapunov exponents are a convenient way to characterize attractors and intercompare different system's attractors.* A good example of this latter point that we make use of in this thesis is relating time scales of disparate models, as discussed at further length in the next chapter.

Much can be said about Lyapunov exponents, and we will not try to summarize all of their interesting aspects. However, we will note that since dissipative systems have contracting areas in their state spaces (i.e., trajectory convergence), their overwhelming average tendency for perturbations made to states on their attractor is to re-attract those perturbed states. That is to say, most Lyapunov exponents are negative — in fact, their sum is always less than zero for a dissipative system (due to Liouville's theorem, e.g., Bergé et al. 1984). Also, for flows, perturbations made along the trajectory path will neither grow nor decay because the dynamics are deterministic and causal. If perturbations are made “into” the attractor (as opposed to “off of” the attractor), thereby displacing a state onto a nearby, neighboring trajectory, then it is possible to have that perturbation grow in amplitude. As we shall see, this is a defining characteristic of a so-called strange attractor. Hence, one can begin to identify chaotic systems by the existence of a positive Lyapunov exponent. This distinguishes chaotic attractor types from other attractor types like so-called *limit cycles*, fixed amplitude oscillations through state space. It should be mentioned that though Lyapunov exponents are often defined with respect to attractors, conservative systems can have Lyapunov exponents as well, only their exponents sum to zero (again, by Liouville's theorem). A positive Lyapunov exponent in that case indicates *nonintegrable* dynamics which are qualitatively similar to chaotic dynamics in that both are sensitive to their initial conditions, while only chaotic systems have attractors.

Before leaving this subject, we briefly contrast Lyapunov exponents with so-called singular values and their corresponding singular vectors: if Lyapunov exponents correspond to an attractor's normal modes, then singular values correspond to a system's state dependent *non-modal* growth (also called non-normal). Lyapunov exponents denote the average growth rate of perturbations over the attractor; singular values characterize a state's realizable perturbation growth factor given its attractor location and a specified time scale of interest. The perturbation growth implied by Lyapunov exponents and singular values can vary quite widely from each other. Similarly, Lyapunov vectors are fixed structures through time, thus allowing only simple amplification; singular vectors can change in time, and thus

are transient in nature. This freedom can allow for localized regions of spectacular perturbation growth (compared to the growth implied by the leading Lyapunov exponent). The concept of localized, state dependent perturbation growth is intimately related to the note we made in our discussion of stability in the previous subsection. If we identify “errors” as perturbations in state space, then DA is intimately related to estimating each component of a state estimate’s error so as to characterize the most likely ensuing (state dependent) error growth — the singular vectors associated with the correct initial error structure (were it known) would give precisely this information.

### 2.2.3 Chaotic systems and strange attractors

As was discussed earlier, the advent of computers allowed not only for prediction but for studies of predictability. Two seminal studies by Edward Lorenz exposed and explained theoretical constraints on certain classes of dynamical systems. The second study examined the consequences of predicting a nonlinear system whose dynamics possesses multiple scales (1969). The conclusion was that uncertainty in scales of motion unresolved by a nonlinear model will cascade upward to effect the largest scales. This is very similar to turbulence ideas, and it presents a hard predictability limit for systems that proceed by approximate numerical methods. This has been expressed popularly as the “butterfly effect,” essentially that a butterfly flapping its wings could generate wind disturbances that are impossible to disprove contributed to the formation of a hurricane or tornado a continent away.<sup>2</sup> This predictability limit is a manifestation of model error, the practical constraint that we will never be able to exactly solve the equations of motion that govern a natural system. As we are mainly concerned with DA, we neglect sources of model error throughout this thesis (there will be discussion of this assumption later).

More germane to our work is Lorenz’s first study where, using only a system of three nonlinear ODEs, Lorenz (1963) explored and explained the evolution of certain systems, termed *chaotic* systems, that can exhibit nonperiodic, transitive, and yet still deterministic solutions. Nonperiodic behavior means that trajectories are not closed circuits through state space, transitive means that there is recurrence to the solutions (i.e., they are confined within a finite hypervolume in state space), yet determinism still constrains trajectories never to intersect. A consequence of this rich dynamical behavior is an extraordinary sensitivity on the initial state, as indicated by the presence of at least one positive Lyapunov

---

<sup>2</sup>Actually, the popular usage of “the butterfly effect” is associated with Lorenz’s first study uncovering the sensitivity to initial conditions, but it is our belief (and also T. Palmer’s (personal communication 2003) and C. Snyder’s (personal communication 2004)) that this second study is the real cause of this phenomenon.

exponent. Sensitivity to initial conditions implies that virtually any two states separated by a vanishingly small difference will evolve to eventually be separated by a much larger, “climatological difference.” All chaotic systems exhibit these two characteristics — sensitivity to initial conditions and having a climate — due to their having *strange attractors*. Strange attractors arise from forced, dissipative dynamics and thus are spatially extended attracting sets through a state space. Trajectories begun on the attractor continue to course through state space bound to the attractor.

### **Fractal dimensions — an attractor’s natural measure**

Geometrically, the conundrum of having nonperiodic, non-intersecting trajectories constrained to a limited manifold of state space can begin to be understood by noting that strange attractors often have fractional dimensions in their systems’ state spaces. Classic examples of geometric creatures with fractional dimensions (so-called “fractals”) are Cantor sets, curves with infinite length contained in a bounded area, or “solids” with zero volume and infinite surface area. The salient feature of objects with fractional dimension is a self-similarity through all scales, where one can see characteristic patterns of organization through many scales of magnification. This self-similarity down through infinitesimal scales allows an attractor to remain bounded without ever intersecting itself. The self-similarity also provides an attractor with infinite “texture,” where identifying a neighborhood of states is relatively easy, but identifying a specific state on the attractor is very difficult. *In fact, if one denotes a given neighborhood on an attractor, then there will in general be an infinite number of non-collinear points in state space occupying that neighborhood.* These states are termed *indistinguishable states*, and their significance will be made clearer soon. The sensitivity to initial conditions follows from a given system evolving every point on its fractal attractor slightly differently and these differences compounding over time.

As an example of this, figure 2-1 shows the *natural measure* of a dynamical system’s attractor. The specific system considered is the Ikeda map, which will be discussed at length in the next chapter, but for now it serves as an example of a generic two-dimensional strange attractor (i.e., the state space is spanned by elements  $x$  and  $y$ ). An attractor’s natural measure is the probability of finding the system’s state in a particular location in state space. The left panel shows the entire attractor, and the right panel shows a magnified region of the left panel centered near  $(0.75, -0.95)$ . To compile this figure we have essentially tallied a high resolution two-dimensional histogram of a very large number of the system’s states. In this manner, we can begin to identify the attractor as the locations in state space that are occupied by the system. The white regions are never occupied, and the colored

regions are colored in proportion to the system’s frequency of visitation (with red being the most frequent). One can begin to sense the self-similarity — the emergence of white spaces of varying sizes interspersed with colored regions. The magnified image on the right indicates similarly proportioned regions of white space to colored space.

*A system’s dynamics set its attractor structure, and therefore it defines the regions it can occupy in state space. The attractor can be thought of as a system’s probability distribution function (PDF).*<sup>3</sup> A PDF is a non-negative, real-valued scalar function whose integral over an  $n$ -dimensional state space is unity. As we shall see, probabilistic DA is a matter of sharpening this initial PDF.

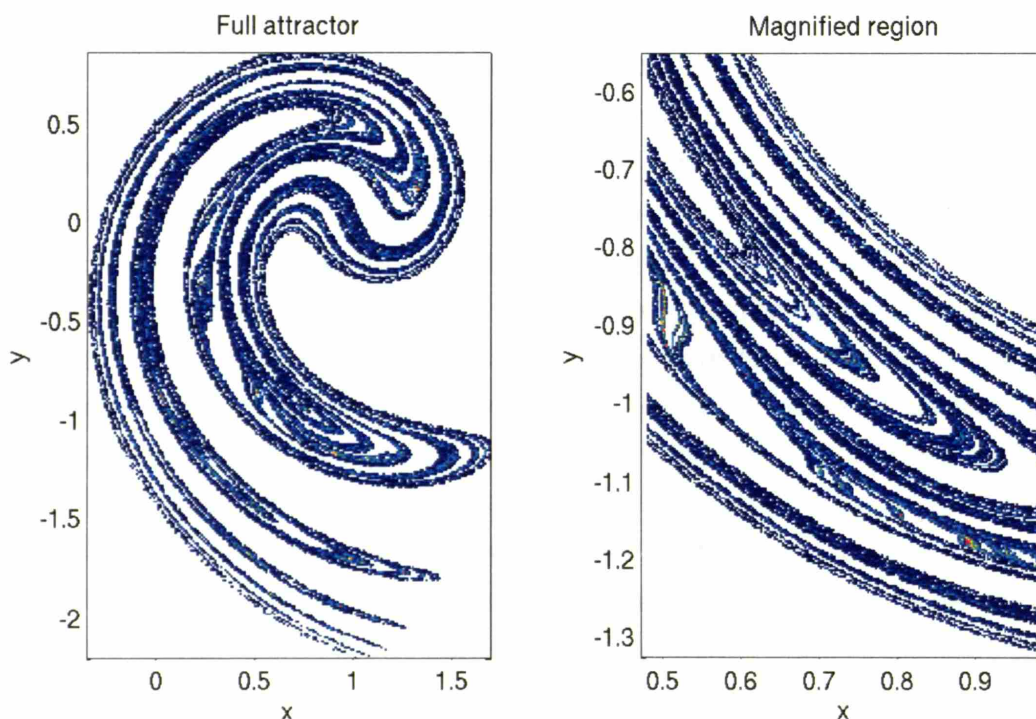


Figure 2-1: These panels show the natural measure of the Ikeda system’s attractor. The left panel shows the full attractor, and the right panel shows a magnified region of the left panel. White spaces indicate regions of state space never occupied by the system, while colored spaces are colored in proportion to the system’s frequency of visitation, with red being the most frequent.

---

<sup>3</sup>Throughout this thesis, “PDF” will stand for both probability distribution function and probability density function, the difference being that only the latter is differentiable. Context should make which is meant apparent.

## 2.2.4 Prediction: geometric interpretation

Simply put, if one has access to a system's attractor geometry and information about its current state, then data assimilation is a matter of locating the current state's location on the attractor, and prediction is a matter of following the current state's trajectory along the attractor. This is what Bjerknes sensed in 1904. However, it is of course not that easy in practice. For physical systems, one never has access to the correct attractor. And model error complications aside, one never has access to the correct state in a chaotic system. Knowledge of the state comes from state *estimation*. If truth is a freely evolving solution of the system, then "error" can be defined by the distance between truth and an estimate of truth in the system's state space (under some suitably defined norm). By chaotic dynamics, the error will tend to grow exponentially. In estimating the true state of a system, one would like to minimize the estimate's error. However, the estimate's error as defined requires knowing truth, which is unknown, so one is forced to employ probability and statistics surrounding the *expected error*.

Judd and Smith (2001) show that *one can never recover the exact true state of a chaotic system even given the correct system dynamics (i.e., no model error) and an arbitrarily long record of noisy observations into the past*. The crux of their argument relies on the existence of the aforementioned indistinguishable states (see section 2.2.3). Essentially, the record of observations helps define a trajectory "tube" through state space, where the cross sectional scale of the tube is set by the observational uncertainty. Judd and Smith showed that one is always able to find "shadowing" trajectories of states distinct from truth that have followed truth through the tube over the observational record. Hence, without precise knowledge (to infinite accuracy!) of the entire true state at one instant in time, one will never be able to identify the true state of the system. This is because no observation is ever errorless. This in itself is a call for *probabilistic* estimation and prediction — the dynamics may be deterministic, but we will never have access to the correct state.

As mentioned at the end of the previous subsection in reference to figure 2-1, a system's attractor can be considered the true state's PDF in state space. Lacking any other information about the state (i.e., a prior state estimate or observations), then the attractor represents an "informed state of ignorance," that is, knowledge of the system dynamics only. Once one begins to have more state-specific information, then one can hone his or her estimate to local neighborhoods of the attractor. The next section addresses methods to achieve this sharpening of uncertainty.

## 2.3 Nonlinear dynamics & predictability

Having reviewed the basics of nonlinear dynamics and the geometric interpretation of DA and prediction, we now review the scientific understanding of predictability that has come to bear since the time of Bjerknes.

### 2.3.1 Prediction: the correct probabilistic approach

In taking a probabilistic approach to state estimation and prediction, one deals with PDFs rather than specific states. Probability is a measure whose global integral is 1 over a space. Some locations in space are more probable than others, hence the notion of a probability density or distribution. A PDF in state space is thus a “cloud” indicating the regions a specific state is likely to occupy. We designate the PDF of a state vector estimate  $\mathbf{x}$  by  $\varphi(\mathbf{x})$ . One may treat an entire PDF, or instead, one often finds it convenient to treat its moments. Central moments are defined about a PDF’s mean. The mean is the expected value of a PDF, defined as:

$$\mu_i = \int x_i \varphi(\mathbf{x}) d\mathbf{x} \equiv \langle x_i \rangle, \quad (2.1)$$

where the subscript  $i$  denotes a particular axis in the  $n$ -dimensional state space. Equation (2.1) defines the expectation operator, which will often be shortened by use of the angle brackets (i.e.,  $\boldsymbol{\mu} = \langle \mathbf{x} \rangle$ ). Note that the expected value is not necessarily the value of maximum likelihood (i.e., the PDF’s mode). As  $\mathbf{x}$  is a state vector, each higher order central moment is a tensor of one higher order. The second central moment is thus a matrix, called the covariance matrix. Its diagonal is the variance of each element. A covariance matrix is non-negative definite and symmetric. Its elements are defined as:

$$P_{ij} = \int (x_i - \mu_i)(x_j - \mu_j) \varphi(\mathbf{x}) d\mathbf{x} = \langle (x_i - \mu_i)(x_j - \mu_j) \rangle. \quad (2.2)$$

Higher order central moments are found similarly.

#### Data assimilation: Bayes rule

Proper probabilistic data assimilation needs to blend the information of a prior PDF state estimate with the PDF of available observations, and the correct way to do this is through Bayes rule (sometimes called Bayes’s theorem). Bayes rule is “correct” by definition, that is, it follows immediately from the definitions of a joint probability density, a conditional probability density, and a marginal probability density.

Bayes rule itself is a straightforward idea; confusion typically arises from its notation. Generically, consider two events,  $A$  and  $B$ . The joint probability density of “ $A$  and  $B$ ,” written  $\varphi(A, B)$ , concerns the likelihood of both events happening in conjunction. If one is interested in the conditional probability density of “ $A$  given  $B$ ,” expressed as  $\varphi(A|B)$ , then one can simply use the definition of the joint probability to evaluate  $\varphi(A|B) = \varphi(A, B)/\varphi(B)$ , where  $\varphi(B)$  is the marginal probability density of event  $B$  alone (from integrating  $\varphi(A, B)$  over  $A$ ). However, if one does not know the joint density, then Bayes rule shows another expression for  $\varphi(A|B)$  is proportional to the conditional probability density of “ $B$  given  $A$ ” times the marginal probability density of  $A$ , expressed simply as  $\varphi(A)$ :

$$\varphi(A|B) \propto \varphi(B|A)\varphi(A). \quad (2.3)$$

The proportionality needs only a normalization factor ( $\varphi(B)^{-1}$ ) to be made an equality (normalization to ensure the probability measure still integrates to 1).

In a prediction context, one is concerned with estimating the state of a system. Hence, define  $\mathbf{x}_t$  as the state of the system at time  $t$ ,  $\mathbf{y}_t^o$  as the observations of the system available at time  $t$ , and  $\mathbf{Y}_t$  as the set of all observations available up to and including time  $t$  (i.e.,  $\mathbf{Y}_t = \{\mathbf{y}_t^o, \mathbf{y}_{t-1}^o, \dots, \mathbf{y}_0^o\}$ ). DA thus seeks the PDF of  $\mathbf{x}_t$  given all the available observations up to and including time  $t$ . However, in a filtering context, one has only a prior estimate of  $\mathbf{x}_t$  at time  $t$  before DA has occurred. The prior estimate is the PDF of  $\mathbf{x}$  conditioned on all observations available before time  $t$ , written as  $\varphi(\mathbf{x}_t|\mathbf{Y}_{t-1})$ . This term corresponds to  $\varphi(A)$  in equation (2.3). The latest observations available constitute new information available about the true state of the system. Their PDF is the conditional probability of having observed  $\mathbf{y}_t^o$  given that the true state is actually  $\mathbf{x}_t$ , and it is written as  $\varphi(\mathbf{y}_t^o|\mathbf{x}_t)$ . This term corresponds to  $\varphi(B|A)$  in equation (2.3). Hence, the Bayesian expression for the probabilistic update of a nonlinear dynamical state estimate is:

$$\varphi(\mathbf{x}_t|\mathbf{Y}_t) = \varphi(\mathbf{y}_t^o|\mathbf{x}_t)\varphi(\mathbf{x}_t|\mathbf{Y}_{t-1})/\text{normalization}. \quad (2.4)$$

This says that the PDF of  $\mathbf{x}_t$  conditioned on all available observations is the normalized product of the observations’ PDF and the prior PDF of  $\mathbf{x}_t$ . *When viewed in a state space context, this is very clear — one is simply multiplying PDFs together.* A good reference for this is D’Agostini (2003).

## One-dimensional Bayes rule example

To further clarify this important point, it is useful to consider a few examples. First consider a one-dimensional example where the state space is a line. Figure 2-2 considers two different prior PDFs and three different observation PDFs. Each panel in the top row has a prior *Gaussian* PDF (also called a normal distribution) of the form:

$$\varphi(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad (2.5)$$

where  $x$  is a point in state space,  $\mu$  is the mean of the PDF, here equal to 0, and  $\sigma^2$  is the variance of the PDF, here equal to 17. This curve is shown in blue. Multi-dimensional extensions exist to equation (2.5) and will be discussed below. A Gaussian PDF is a special PDF because it is completely specified by its first two central moments: all higher odd moments are 0 and all higher even moments are determined by the second moment. From information theory, this means that of all PDFs with the same specified first two central moments, the Gaussian PDF has the lowest information content (i.e., maximal “spreading”) (Tarantola 1987). A Gaussian PDF is often the easiest and most familiar PDF with which to work.

In contrast to the Gaussian PDF, the bottom row of panels in figure 2-2 each has a prior bimodal PDF of the form:

$$\varphi(x) = \frac{1}{2\sqrt{2\pi}} \left[ \exp\left(-\frac{1}{2}(x-\nu)^2\right) + \exp\left(-\frac{1}{2}(x+\nu)^2\right) \right], \quad (2.6)$$

where  $\nu$  is the offset of each peak from 0, here equal to 4. It turns out that this bimodal PDF has the exact same mean and variance as the Gaussian PDF described in equation (2.5), however, unlike the Gaussian, it takes more information than a mean and variance to define the bimodal PDF. This curve is also shown in blue.

Each of the six panels shows an update of one of the two prior PDFs by a second PDF, representing a newly available observation. The observation PDFs are Gaussian PDFs as well (see equation (2.5)), but they each have  $\mu = 3.5$ , a value within one standard deviation of the prior Gaussian’s mean and also consistent with a true state somewhere within the positive peak of the prior bimodal PDF. The three considered observation PDFs differ in their variances (i.e., their accuracy). The first column of panels in figure 2-2 considers a rather accurate observation with a variance one quarter that of the prior PDFs’. The second column uses an observation PDF with a variance equal to the prior PDFs’, and the third column uses a rather inaccurate observation with a PDF that has four times the prior PDFs’

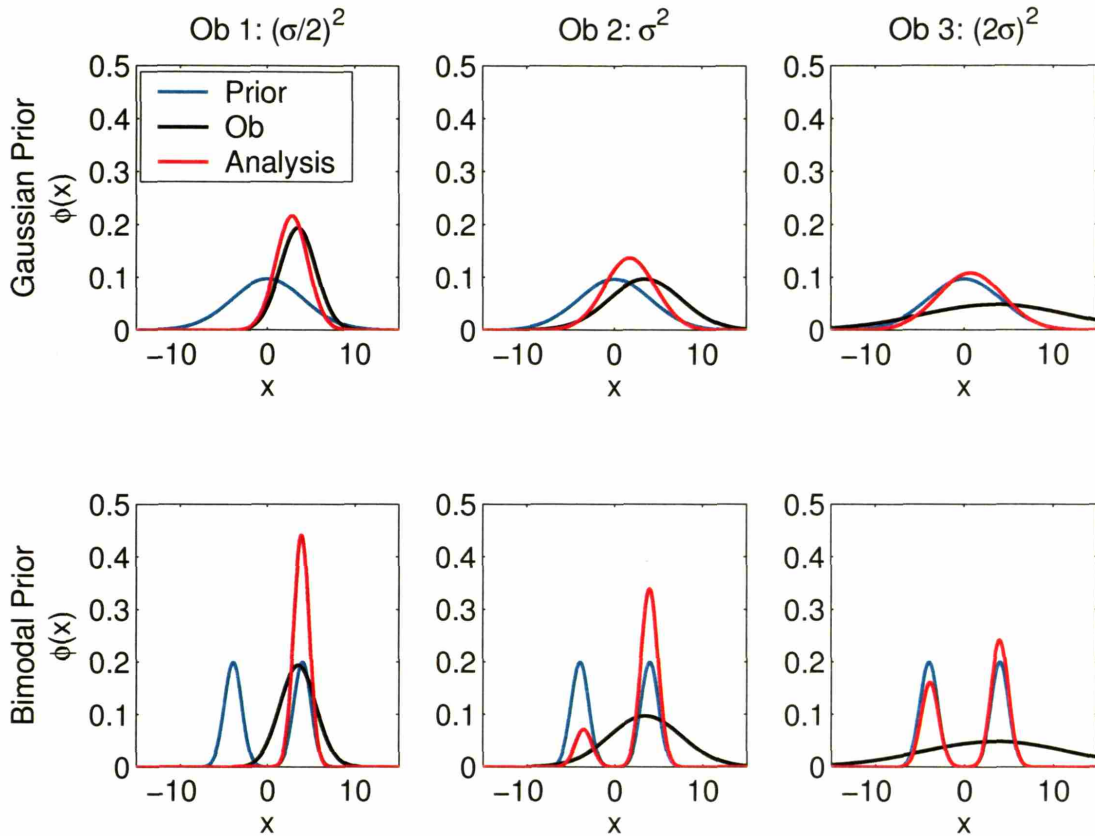


Figure 2-2: One dimensional example of Bayes rule updates. The top row of panels each assume a Gaussian prior, and the bottom row assumes a bimodal prior. Each column uses an observation PDF with a different accuracy,  $(\sigma/2)^2$ ,  $\sigma^2$ , and  $(2\sigma)^2$ , respectively.

variance. The observation PDF curves are all drawn in black. Each panel also shows a red curve, the result of evaluating Bayes rule using the prior and observation in that panel. *It is clear that the obtained **analyses** (the updated PDFs) are simply the normalized product of their prior and the observation.* For the Gaussian prior, the analysis resembles the more accurate information source, as common sense would dictate. For the bimodal prior, the two initially equal peaks come to receive unequal weights depending on the accuracy of the observation. Also, all of the analyses retain the prior knowledge that there is essentially zero probability of having a state at  $x = 0$ . Bayes rule accomplishes just what one would expect.

## Two-dimensional Bayes rule example

As all of the dynamical systems considered in this thesis are multi-dimensional, it is also instructive to consider a two-dimensional example of a Bayes rule update. In this example, the prior is a 2D Gaussian distribution. The  $n$ -dimensional extension to equation (2.5) is as follows:

$$\varphi(\mathbf{x}) = ((2\pi)^n \det(\mathbf{P}))^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{P}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right), \quad (2.7)$$

where  $\boldsymbol{\mu}$  is the mean vector as defined in equation (2.1),  $\mathbf{P}$  is the covariance matrix as defined in equation (2.2), and  $\det(\cdot)$  is the matrix determinant. A multi-dimensional Gaussian in state space is an  $n$ -dimensional hyperellipse; the eigenvectors of  $\mathbf{P}$  determine the relative ranking of major and minor axes. In two dimensions, Gaussians are simple ellipses (circles if  $\mathbf{P}$  is proportional to the identity matrix). In the case where the state is fully observed, both the prior PDF and the observations' PDF are ellipses. Bayes rule operates precisely the same way as on the one-dimensional line. The results are shown in figure 2-3. The large square panel shows a neighborhood of the 2D state space. There are three PDFs contoured in the plot: the blue contours are the elliptic Gaussian prior PDF, the black contours are the observations' circular PDF, and the red contours are the Bayes rule analysis from the other two. Each PDF is contoured with increments of its standard error, that is, the square-root of the variance in each axis direction of the ellipses. *Again, the analysis is clearly just the normalized product of the other two PDFs.* The two side panels show the marginal PDFs for each of the two state variables. Strictly, Bayes rule must be applied in the full state space rather than on each individual marginal PDF, but the marginals emphasize how the PDF was made sharper.

If the state were not fully observed, then the clearest geometric interpretation is to simply assume a state of total ignorance in the unobserved directions. Most of the time, this will be a uniform PDF, but strictly it depends on the norm and coordinates of the state space (Tarantola 1987). This is completely fine because if one finds that if non-zero covariance information exists between state elements (i.e., non-zero off-diagonal terms in  $\mathbf{P}$ ), then knowledge of one element has predictive power over others. In conclusion, we emphasize a point made by Tarantola (1987): Bayes rule is a mathematical tautology which only becomes physically applicable once postulates have attached concrete interpretations to the involved PDFs. Stated another way, *the algebra of PDFs is independent of the interpretation of PDFs.* However, here in the process of filtering the state of a nonlinear dynamical system, the interpretations are clear, and the approach bears little controversy.

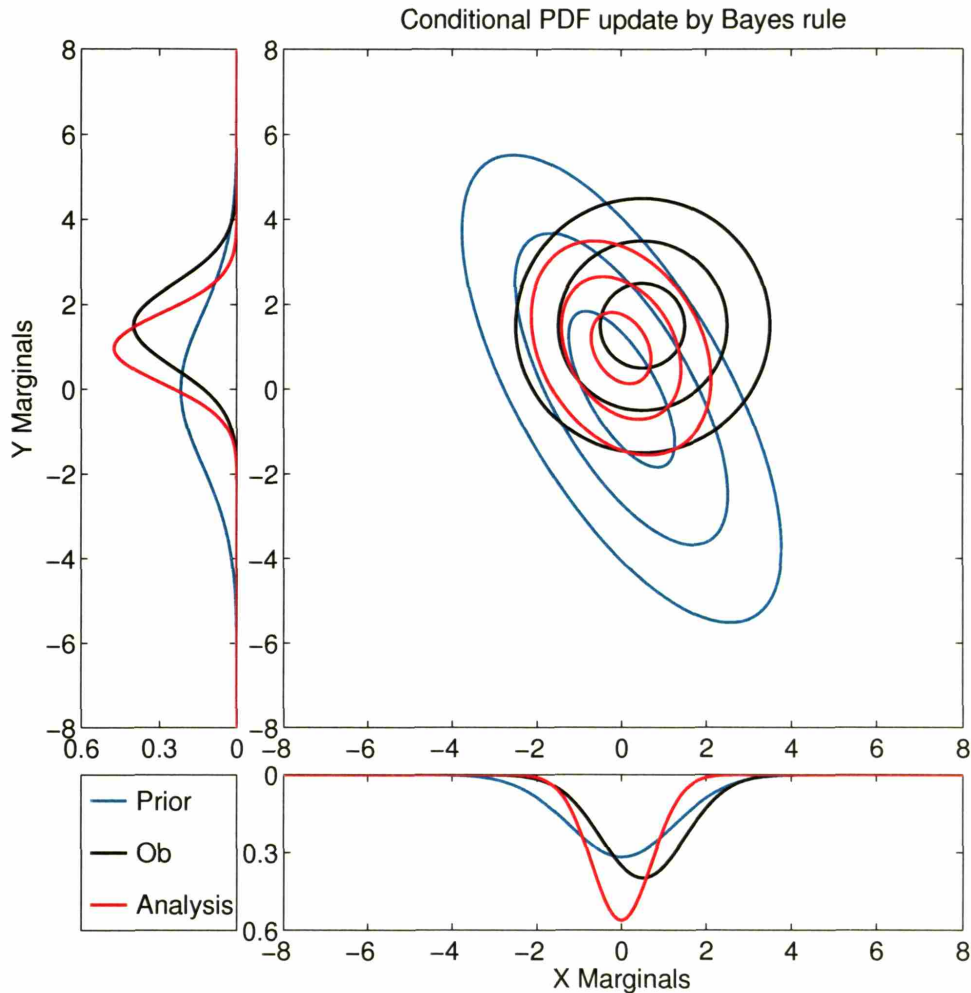


Figure 2-3: Two-dimensional example of Bayes rule updates. The large panel shows contours of three 2D Gaussian PDFs. The contour interval for each PDF is in terms of its axes' standard deviations. The side panels show the marginal PDFs for each state variable.

### Forecasting: Fokker-Planck equation

State estimation is only one half of the filtering process; the other half is the evolution of the state estimate using the system dynamics. Use of the system dynamics is implied in our interpretation of the prior PDF in equation (2.4) in the Bayes rule discussion above. DA at time  $t - 1$  produced the best estimate  $\varphi(\mathbf{x}_{t-1}|\mathbf{Y}_{t-1})$ , and invoking it as the prior estimate at time  $t$  implies that a model was used to evolve the estimate from  $t - 1$  to  $t$  so that it is valid at time  $t$ . And in fact, that is what the system dynamics do, but it is worth questioning how a deterministic system can evolve a PDF.

Generally, from a deterministic point of view, the system dynamics, including all its

parameters, can be represented by a nonlinear operator that evolves a state forward in time. In continuous time, one would consider  $d\mathbf{x} = \mathcal{M}^c dt$ , where  $\mathcal{M}^c$  is the time-continuous system dynamics, but as most numerical methods require use of a timestep, one can generically represent the timestepping routine as well:

$$\mathbf{x}_{i+1} = \mathcal{M}[\mathbf{x}_i], \quad (2.8)$$

where  $\mathcal{M}[\cdot]$  is a nonlinear operator that advances the state one timestep.

To generalize the deterministic evolution of a state to the evolution of a PDF, one can use the Fokker-Planck equation (FPE), also called the forward Kolmogorov equation. The FPE is the cornerstone of evolving stochastic dynamical systems, and in its full complexity includes terms for noise induced drift and probability diffusion (Penland 2003). However, when applied to deterministic dynamics, the full equation reduces to the Liouville equation, which for a PDF is nothing more than a continuity equation for probability in a system's  $n$ -dimensional state space:

$$\frac{\partial \varphi(\mathbf{x}, t)}{\partial t} + \nabla_n \cdot (\dot{\mathbf{x}} \varphi(\mathbf{x}, t)) = 0, \quad (2.9)$$

where  $\nabla_n$  is an  $n$ -dimensional gradient operator, and  $\dot{\mathbf{x}}$  is the time rate of change at a point,  $\mathbf{x}$ , in state space. This term is given by the system dynamics, and can be visualized as providing a “probability current” through state space.

### 2.3.2 Approximate probabilistic approaches to prediction

As discussed in the previous subsection, subsection 2.3.1, if one is performing strict probabilistic prediction of a nonlinear dynamical system, one should be using Bayes rule for DA and the FPE for forecasting. However, the main problem with the correct approach is that it requires specification of the state's PDF. In NWP, a single state can require more than  $10^7$  numbers; a PDF is an extended creature through state space (except for  $\delta$ -function PDFs) whose representation would properly require *discretizing state space* itself. If one were to resolve  $10^2$  values per orthogonal axis, then the number of state space gridpoints required to represent  $\varphi$  would be  $(10^7)^{10^2} = 10^{700}$  (!) Solving equation (2.9) is prohibitively expensive for modestly sized problems as well. From an information point of view, such a multitude of numbers may well be required to properly do probabilistic prediction; however, like most unapproachable scientific pursuits, one makes progress through careful approximation. *There are alternative approaches to prediction that reach equivalence with the*

*correct method in the limit of infinite resources (e.g., information and computation).* These alternative approaches in truncated form may be easier to implement: one approach is to consider the moments of a PDF and the other approach is to approximate the PDF by Monte Carlo methods. These alternative methods lead naturally to approximations of the correct approach.

### **PDF moment closure methods**

Knowledge of all the moments of a PDF is equivalent to knowledge of the PDF itself, however, as noted by Epstein (1969), complete evaluation of a PDF “both requires and provides more information than is normally of interest” (page 742). To capitalize on this observation, Epstein reframed the FPE as individual evolution equations for a PDF’s moments. In his *stochastic dynamic prediction* approach, he found that evolving a given moment of a PDF requires knowledge of the next higher moment (e.g., the mean requires knowledge of the covariance; this phenomenon is quite like turbulent cascades through length scales in fluids). *Hence, in order to exactly duplicate the results of the FPE, one must evolve an infinite number of moments; however, in order to approximate the FPE, one need only invoke a moment closure scheme.* Epstein suggested a simple Gaussian closure scheme (i.e., third moment discard), but any number of elaborate closure schemes is possible. An excellent review of this material is given in Cohn (1993). The complementary moment-based DA scheme to stochastic dynamic prediction is still Bayes rule, only tapered to fit the particular moment closure assumption. For a Gaussian closure, this reduces to the method of least squares weighted by the PDFs’ covariances, or equivalently, the Kalman filter update. Least squares and the Kalman filter will be treated in more detail below.

### **Monte Carlo methods**

Monte Carlo methods provide another approach to solving the FPE. Monte Carlo methods are based on explicit population of a PDF’s “cloud” of uncertainty within state space by  $N$  random realizations from it. Each state is then explicitly evolved forward under the system dynamics. The states’ normalized aggregate evolution will approximate the evolution of their original PDF. The populating states are often called ensemble members, particles, or replicates. Monte Carlo methods will converge to the evolution of  $\varphi$  governed by the FPE (equation (2.9)) as the number of particles tends to infinity; however, the convergence is rather slow, as  $N^{-1/2}$ . *If a finite number of particles are used, then Monte Carlo methods provide a convenient approximation to the correct approach.* Lorenz (1965) and Leith (1974) were among the first to suggest and use Monte Carlo methods in atmospheric science, though

Monte Carlo methods themselves have been in use much longer, even back to the late 1700s (e.g., Wikipedia 2001).

### Particle filters

Monte Carlo methods' complementary DA approach is again based on Bayes rule, as would be expected. However, as Bayes rule multiplies PDFs, it needs to be modified to handle particle realizations from those PDFs. The resolution to this is typically called a *particle filter* or a weighted ensemble filter. As cited by Kim et al. (2003), the particle filter was apparently originally formulated during the Manhattan Project by Ulam and von Neumann, though was not popularly accepted until the work of Enăchescu (1985) and later Gordon et al. (1993). *The key aspect of the particle filter is not to modify any of the particle states themselves, but to modify the weight each state is given in their aggregate.* To this, initially each particle has equal weight  $w_i = \frac{1}{N}$ , where  $i = 1, N$ . The weights remained fixed during model evolution. During DA, available observations are allowed to update the particle weights based on the particles' relative distances from the observations.

The easiest implementation assumes that the observations have Gaussian PDFs. Assume that the observations available at time  $t$  are stored in an  $m$ -dimensional column vector  $\mathbf{y}^o$  with error covariance matrix  $\mathbf{R}$ . If  $m \neq n$ , or if the observations are indirect (i.e., not of elements in the state vector), then one must invoke a possibly nonlinear operator to translate between “observation space” and state space. Here we generically represent all such transformations and interpolations by the operator  $\mathcal{H}[\cdot]$ . Use of this operator (along with a chosen norm) allows one to find the distances between the particles and the observations in observation space. These distances then inform the reweighting of the particles:

$$w_i^+ = \frac{w_i}{\text{normalization}} \exp \left( -\frac{1}{2} (\mathbf{y}^o - \mathcal{H}[\mathbf{x}_i])^T \mathbf{R}^{-1} (\mathbf{y}^o - \mathcal{H}[\mathbf{x}_i]) \right), \quad (2.10)$$

where  $w_i^+$  are the updated weights, and the normalization serves to ensure that  $\sum_{i=1}^N w_i^+ = 1$  remains true after the updates. The argument of the exponential is a quadratic form, quite like that shown in the definition of a multi-dimensional Gaussian (see equation (2.7)). It weights the actual distance between the observations and the state,  $\mathbf{y}^o - \mathcal{H}[\mathbf{x}]$ , termed the *innovation*, by the expected distance as stored within the covariance matrix. The smaller the innovation, the larger  $w_i$  becomes.

One of the main criticisms of the particle filter is the so-called “curse of dimensionality,” the exponential growth of hypervolume as a function of dimension. *As  $n$  increases, the*

*number of particle realizations necessary to maintain a given level of resolved accuracy increases exponentially.* This gives rise to the “empty space phenomenon,” where the tails of high-dimensional PDFs become quite important. Silverman (1986) points out that even for a ten-dimensional Gaussian PDF with zero mean and unit variance in each dimension, 99% of the density’s mass is at points greater than 1.6 away from the mean, quite different than the one and two-dimensional Gaussians considered above. Anderson and Anderson (1999) discuss how this can weaken particle filters that update weights.

Another problem is that since the particles themselves are never altered, only their weights, it would seem that a nearly infinite number of particles is required to track the true trajectory of a chaotic system since most particles will come to be given zero weighting as they exponentially diverge from the observations. Only those shadowing trajectories that trace through the observation tube referred to in section 2.2.4 will survive a lengthy filtering exercise with a particle filter. Hence, the particle filter is most often presented as a weighted *resampling* filter, where the particle distribution is actually resampled from the approximate PDF after each update for their weights. One of the attractive aspects of the particle filter is that the states are never modified, and hence they are able to stay on the system’s attractor (or at least they can be given a chance to find it if they were initially realized off of it). If one resamples from the approximate PDF, there is no guarantee that the newly realized states will be on the attractor. Therefore, one often resamples from the states represented before DA was begun (i.e., the prior ensemble); new versions of the old states are drawn in proportion to their newly devised weights. This has the consequence of possibly forming several replicates of the exact same state. If the dynamics are deterministic, then all duplicate states will evolve the same and important degrees of freedom will have been lost: indeed, after many successive updates, one can imagine all particles collapsing onto one state. Several approaches have been devised to combat this: so-called kernel density approaches (Anderson and Anderson 1999), including stochastic system noise (van Leeuwen 2003), and performing some kind of deterministic perturbing of states (Kim et al. 2003). All proposed methods have drawbacks, and we note that resampling filters are also subject to the curse of dimensionality.

### **2.3.3 Practical probabilistic approaches to prediction**

It is with some frustration that much from the beautiful theory from nonlinear dynamics and state space considerations cannot actually be implemented in practical problems — a correct approach is known, as are at least two plausible approximations thereof, yet none of these is implementable in problems of practical interest. To implement these methods,

there are two major obstacles to surmount: 1) *The first obstacle is an obvious computational limitation.* As resolution becomes finer and state vectors' dimensions grow larger, so does the computational burden of evolving PDFs and their moments. Yet Moore's law persists at the same time humans are making large strides in distributed computing. It is not inconceivable that computational burdens will someday disappear; however, if history has taught a lesson, it is likely that scientists will continue to pour added computational resources into increased system resolution rather than further advances in evolving and estimating PDFs. 2) *The second important problem facing scientists is the lack of available information to properly close prediction problems.* Dee (1991) argues strongly that this is the real implementation obstacle to surmount; he feels scientists may never have enough information to specify even very simple schemes that consider only the mean and covariance of the involved PDFs, let alone more complex and complete schemes. It is apparent that Monte Carlo methods attempt to sidestep this information paucity through increased computational demand as the full covariance information is approximated, but one still need know how to sensibly initialize the particles. Hence, because of these obstacles and other operational considerations, scientists have been forced to make even further approximations to the correct method.

### **Gaussian error PDFs & linear error growth**

The most common and pervasive assumptions made throughout the literature and the community is to assume that the dynamics are linear, or can at least be validly linearized over time periods of interest, and that the involved PDFs are strictly Gaussian. These two assumptions conspire to require that errors remain in regimes of linear error growth, and they are related because nonlinear dynamics can readily evolve a Gaussian PDF into a non-Gaussian one. One attractive aspect of these assumptions is that prediction and state estimation begin to more closely resemble the initial (naïve) deterministic approaches to prediction used before much was known about chaotic, nonlinear dynamical systems: the mean of the Gaussian PDF is *the* state estimate and the PDF's covariance is the state estimate's uncertainty. This makes presentation and communication of predictions easier. Also, it allows for deterministic forecasting verification techniques. Another attractive aspect of these assumptions is that, if met, they allow expression of various "optimal" state estimates, where optimal is a user-defined notion. One common condition of optimality is that the state estimate have minimum variance compared to all others based on the information at hand. Another is that the state estimate be the one of maximum likelihood (i.e., the PDF's mode). As it turns out, having linear dynamics ensures that the obtained

estimate is the minimum variance estimate (i.e., the PDF's mean), and having Gaussian PDFs further ensures that the obtained estimate is the maximum likelihood estimate (e.g., Jazwinski 1970).

### Least squares

Once one has assumed linear(izable) dynamics and Gaussian uncertainties, then Bayes rule can be shown to reduce to a specially weighted form of linear least squares estimation. Devising the method of least squares is perhaps the first attempt at DA, and is often attributed to Gauss as he sought to predict the motions of the planets, though some cite Legendre for its development (both were in the early 1800s). The basic notion is to minimize the sum of the squared distances of an estimate to its contributing information sources. Distances are computed under an assumed norm, and least squares assumes the  $\ell^2$  norm (i.e., simple Euclidean distance). When minimizing the distances, one can weight them with respect to one another to give more credibility to one information source over another. Bayes rule reduces to least squares when the weights used are the PDFs' covariances. If both the prior PDF and the observations' PDF are Gaussians, then their product is the sum their exponents. Referring back to equation (2.7):

$$\begin{aligned}\varphi^a &\propto \varphi^f \cdot \varphi^o \\ \varphi^a &\propto \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}^f)^T (\mathbf{P}^f)^{-1} (\mathbf{x} - \mathbf{x}^f) - \frac{1}{2}(\mathcal{H}[\mathbf{x}] - \mathbf{y}^o)^T \mathbf{R}^{-1} (\mathcal{H}[\mathbf{x}] - \mathbf{y}^o)\right),\end{aligned}\tag{2.11}$$

where  $\varphi^a$  is the updated PDF, termed the *analysis* PDF,  $\varphi^f$  is the prior PDF,  $\varphi^o$  is the observations' PDF.  $\varphi^f$  is assumed to be Gaussian with mean  $\mathbf{x}^f$  and covariance  $\mathbf{P}^f$ , and  $\varphi^o$  is assumed to have mean  $\mathbf{y}^o$  and covariance  $\mathbf{R}$ . Throughout this thesis we have tried to stick with the DA notation established by Ide et al. (1997). Equation (2.11) shows that  $\varphi^a$  will also be Gaussian, and its mean and covariance are determined by  $\mathbf{x}^f$ ,  $\mathbf{P}^f$ ,  $\mathbf{y}^o$ , and  $\mathbf{R}$ . The mean of a Gaussian is also its mode, meaning that one can find it by finding the value of  $\mathbf{x}$  that maximizes  $\varphi^a$ . Maximizing  $\varphi^a$  is equivalent to minimizing  $-\log(\varphi^a)$ . The exponent in equation (2.11) then defines a cost function, or penalty function,  $J$ , that is the sum of the weighted distances between the mean of  $\varphi^a$  (i.e., the sought state estimate) and the various available information sources. To achieve the strict optimality results, the operator  $\mathcal{H}[\cdot]$  must be linear. The operator is thus written as a matrix,  $\mathbf{H}$ , to emphasize its linearity:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x}^f - \mathbf{x})^T (\mathbf{P}^f)^{-1} (\mathbf{x}^f - \mathbf{x}) + \frac{1}{2}(\mathbf{y}^o - \mathbf{H}\mathbf{x})^T \mathbf{R}^{-1} (\mathbf{y}^o - \mathbf{H}\mathbf{x}),\tag{2.12}$$

where the ordering of vector differences has been reversed to comply with least squares convention.<sup>4</sup> In general, the method of least squares seeks the minimum of a cost function. That Bayes rule reduces to a quadratic cost function when both  $\varphi^f$  and  $\varphi^o$  are Gaussian PDFs is somewhat of a coincidence, allowing it to be identified as a specially weighted least squares problem.

There are many solution methods one might apply to equation (2.12), and as mentioned in 2.1.1, they are broadly grouped under so-called variational methods and sequential methods. In general, variational methods approximate the solution to equation (2.12) by iterative methods, often informed by the cost function's gradient, whereas sequential methods attempt to directly solve for the state that minimizes equation (2.12). For instance:  $J$  is a scalar cost, but its derivative with respect to  $\mathbf{x}$  is a vector:

$$\frac{dJ}{d\mathbf{x}} = -(\mathbf{P}^f)^{-1} (\mathbf{x}^f - \mathbf{x}) - \mathbf{H}^T \mathbf{R}^{-1} (\mathbf{y}^o - \mathbf{H}\mathbf{x}), \quad (2.13)$$

and the desired mean vector,  $\mathbf{x}^a$ , is the solution to equation (2.13) when  $\frac{dJ}{d\mathbf{x}} = 0$ . Solving this equation with the use of some matrix inversion identities yields the expression:

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} (\mathbf{y}^o - \mathbf{H}\mathbf{x}^f). \quad (2.14)$$

This is the expression for the mean of  $\varphi^a$ , and a similar expression can be found for its covariance. To find this, one notes that covariance matrices are added "in parallel," that is their inverses are additive (the so-called information form of covariances). To find this from equation (2.12), one first finds the Hessian matrix of  $J$ :

$$(\mathbf{P}^a)^{-1} = \frac{d^2 J}{d\mathbf{x}^2} = (\mathbf{P}^f)^{-1} + \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H}.$$

To find  $\mathbf{P}^a$  itself, one need only invert the RHS of the above expression. With the use of some more matrix inversion identities, one can find:

$$\mathbf{P}^a = \mathbf{P}^f - \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H} \mathbf{P}^f. \quad (2.15)$$

Both equations (2.14) and (2.15) can be simplified by defining a matrix

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}. \quad (2.16)$$

---

<sup>4</sup>Quadratic forms are non-negative definite so the ordering does not matter as long as all terms are consistent.

Use of  $\mathbf{K}$  enables the expressions for the mean and covariance of  $\varphi^a$  to be written:

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{K} (\mathbf{y}^o - \mathbf{H}\mathbf{x}^f) \quad (2.17)$$

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{K}\mathbf{H}) \mathbf{P}^f, \quad (2.18)$$

where  $\mathbf{I}$  is an  $n$ -dimensional identity matrix. Examination of these equations reveals much. First, noting that  $\mathbf{H}$  exists to translate between the  $n$ -dimensional state and the  $m$ -dimensional observations, ignoring it shows that the matrix  $\mathbf{K}$  is simply comparing the expected uncertainties,  $\mathbf{K} \sim \mathbf{P}^f / (\mathbf{P}^f + \mathbf{R})$ . Because  $\mathbf{P}^f$  and  $\mathbf{R}$  are both covariance matrices,  $\mathbf{K}$  is therefore a non-negative definite weighting matrix that ranges between a matrix of zeros when  $\mathbf{R} \gg \mathbf{P}^f$  and the identity matrix when  $\mathbf{R} \ll \mathbf{P}^f$ . In this optimal setting,  $\mathbf{K}$  is often called the Kalman gain matrix for reasons to be explained later.

Equation (2.17) shows that the analysis mean is equal to the prior mean plus an additive correction, termed the analysis *increment* (i.e.,  $\mathbf{x}^a = \mathbf{x}^f + \Delta\mathbf{x}$ ). The increment is the matrix product of the gain matrix and the innovation vector (i.e., the  $m$ -dimensional difference between the observations and the prior mean). An additive analysis increment means that the analysis mean is a linear combination of the prior mean and the observations. Geometrically within state space, this is moving the estimate in a straight line direction to compensate for the expected error. The increment will be made zero (i.e., analysis is equal to the prior) if either the innovation is zero, meaning that there are no discrepancies between the available information sources, or if the gain matrix is zero, meaning that the prior mean is very much more accurate than the observations. The increment will replace the prior with observations when the gain matrix is the identity. *Hence, the analysis is simply a weighted linear combination of the prior estimate and the observations.* The analysis is optimal in both a maximum likelihood sense (by virtue of the derivation) and a minimum variance sense (by virtue of the fact that the variance is always minimized about a PDF's mean).

Equation (2.18) shows that the analysis covariance matrix is equal to the prior covariance matrix less a non-negative definite correction. *This means that the analysis covariance is always reduced from, or at worst kept equal to, the prior — essentially, this is an assertion that added information from observations can only serve to sharpen an estimate's accuracy.*

## The Kalman filter

Equations (2.17) and (2.18) have been derived from Bayes rule by assuming Gaussian PDFs for both the prior and the observations. It turns out that it is equivalent to using least squares weighted by the PDFs' covariances. As stressed by Jazwinski (1970), this is some-

what of a coincidence as least squares has no formal probabilistic interpretation, it is merely trying to minimize misfits between the analysis and other information sources. On the other hand, interpretation of the prior and observations as random realizations from PDFs suggests the Bayesian idea of maximizing the conditional PDF  $\varphi(\mathbf{x}_t|\mathbf{Y}_t)$  with respect to  $\mathbf{x}_t$ . Wunsch (1996) also stresses the coincidence that least squares agrees with the optimal approach. Another approach to obtaining the same equations is via the Kalman filter (e.g., Jazwinski 1970). There are many possible derivations of the Kalman filter, including Kalman's original method of orthogonal projections (Kalman 1960), seeking a maximum likelihood estimate (e.g., Lorenc 1986), seeking a minimum variance estimate (e.g., Cohn 1997), or simply identifying it as a specially weighted version of recursive least squares (Wunsch 1996). As the Kalman filter and its variations are the most heavily considered filtering scheme in this thesis, we present a derivation here. This derivation seeks a minimum variance estimate, but is more general than most as it allows for correlations between the prior estimate and the observations. It is adapted from Liebelt (1967).

### The Kalman filter update step

Denote the true state of a system as  $\mathbf{x}^t$ . The prior PDF estimate is a Gaussian with a mean vector  $\mathbf{x}^f$  and a covariance matrix  $\mathbf{P}^f$ , and similarly, the observations' PDF is a Gaussian with a mean vector  $\mathbf{y}^o$  and covariance matrix  $\mathbf{R}$ . One seeks the analysis PDF, or equivalently, its mean vector,  $\mathbf{x}^a$ , and its covariance matrix,  $\mathbf{P}^a$ . Consider the error vectors separating the PDFs' means from the true state:

$$\mathbf{x}^t = \mathbf{x}^f + \boldsymbol{\varepsilon}^f \tag{2.19}$$

$$\mathbf{H}\mathbf{x}^t = \mathbf{y}^o + \boldsymbol{\varepsilon}^o \tag{2.20}$$

$$\mathbf{x}^t = \mathbf{x}^a + \boldsymbol{\varepsilon}^a. \tag{2.21}$$

If the error vectors are unbiased, then:

$$\begin{aligned} \langle \boldsymbol{\varepsilon}^f \rangle &= \mathbf{0}, & \langle \boldsymbol{\varepsilon}^o \rangle &= \mathbf{0}, & \langle \boldsymbol{\varepsilon}^a \rangle &= \mathbf{0} \\ \langle \boldsymbol{\varepsilon}^f \boldsymbol{\varepsilon}^{fT} \rangle &= \mathbf{P}^f, & \langle \boldsymbol{\varepsilon}^o \boldsymbol{\varepsilon}^{oT} \rangle &= \mathbf{R}, & \langle \boldsymbol{\varepsilon}^a \boldsymbol{\varepsilon}^{aT} \rangle &= \mathbf{P}^a, \end{aligned}$$

where  $\mathbf{0}$  is an appropriately dimensioned column vector of zeros. Typically, one assumes that the expected errors in the prior PDF are uncorrelated to the expected errors in the

observations' PDF, but here we allow for this possibility:

$$\mathbf{G} \equiv \langle \boldsymbol{\varepsilon}^f \boldsymbol{\varepsilon}^{oT} \rangle, \quad (2.22)$$

so that  $\mathbf{G}^T = \langle \boldsymbol{\varepsilon}^o \boldsymbol{\varepsilon}^{fT} \rangle$ . The covariance matrices  $\mathbf{P}^f$ ,  $\mathbf{R}$ , and  $\mathbf{P}^a$  are square and symmetric, whereas  $\mathbf{G}$  need not be (it is  $n \times m$ ).

To derive the Kalman filter begin by assuming the analysis PDF mean can be written as a linear combination of the prior PDF mean and the observations. As before, the matrix  $\mathbf{H}$  serves to transform from  $n$ -dimensional state space to  $m$ -dimensional observation space. Hence, with the benefit of foresight, assume the linear combination is written as:

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{K} (\mathbf{y}^o - \mathbf{H}\mathbf{x}^f), \quad (2.23)$$

where the weights of the linear combination are in the  $n \times m$  matrix  $\mathbf{K}$ . Since equation (2.23) is linear, one is free to subtract  $\mathbf{x}^t$  from both sides and to add and subtract  $\mathbf{H}\mathbf{x}^t$  within the parantheses to arrive at an equation for the errors:

$$\boldsymbol{\varepsilon}^a = \boldsymbol{\varepsilon}^f + \mathbf{K} (\boldsymbol{\varepsilon}^o - \mathbf{H}\boldsymbol{\varepsilon}^f). \quad (2.24)$$

The minimum variance analysis PDF will have the minimum diagonal values of its covariance matrix. Therefore, we seek the matrix  $\mathbf{K}$  that minimizes  $\text{diag}(\mathbf{P}^a) = \text{diag}(\langle \boldsymbol{\varepsilon}^a \boldsymbol{\varepsilon}^{aT} \rangle)$ . To find this, evaluate the outer product of each side of equation (2.24) with itself, and then apply the expectation operator. Doing this and rearranging a bit yields the expression:

$$\mathbf{P}^a = \mathbf{P}^f + \mathbf{K} (\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R} - \mathbf{H}\mathbf{G} - \mathbf{G}^T\mathbf{H}^T) \mathbf{K}^T - (\mathbf{P}^f\mathbf{H}^T - \mathbf{G}) \mathbf{K}^T - \mathbf{K} (\mathbf{H}\mathbf{P}^f - \mathbf{G}^T). \quad (2.25)$$

This expression allows use of a matrix algebra generalization of completing the square:

$$\mathbf{A}\mathbf{C}\mathbf{A}^T - \mathbf{B}\mathbf{A}^T - \mathbf{A}\mathbf{B}^T = (\mathbf{A} - \mathbf{B}\mathbf{C}^{-1}) \mathbf{C} (\mathbf{A}^T - \mathbf{C}^{-T}\mathbf{B}^T) - \mathbf{B}\mathbf{C}^{-1}\mathbf{B}^T,$$

where  $\mathbf{A} = \mathbf{K}$ ,  $\mathbf{B} = \mathbf{P}^f\mathbf{H}^T - \mathbf{G}$ , and  $\mathbf{C} = \mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R} - \mathbf{H}\mathbf{G} - \mathbf{G}^T\mathbf{H}^T$ . Completing the square of the RHS of equation (2.25) yields (in mixed notation):

$$\mathbf{P}^a = \mathbf{P}^f + (\mathbf{A} - \mathbf{B}\mathbf{C}^{-1}) \mathbf{C} (\mathbf{A}^T - \mathbf{C}^{-T}\mathbf{B}^T) - \mathbf{B}\mathbf{C}^{-1}\mathbf{B}^T.$$

Because covariance matrices and their inverses are positive definite, all three matrix terms on the RHS of the above equation have positive diagonal values. Therefore, the way to

minimize  $\text{diag}(\mathbf{P}^a)$  is to zero out the second term's contribution, that is have  $\mathbf{A} - \mathbf{B}\mathbf{C}^{-1} = \underline{\mathbf{0}}$ , where  $\underline{\mathbf{0}}$  is an appropriately dimensioned matrix of zeros. By the definitions of  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$ , this implies the sought variance minimizing matrix is:

$$\mathbf{K} = (\mathbf{P}^f \mathbf{H}^T - \mathbf{G}) (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R} - \mathbf{H} \mathbf{G} - \mathbf{G}^T \mathbf{H}^T)^{-1}. \quad (2.26)$$

This is the Kalman gain matrix, and it weights a linear combination of the means of two PDFs in such a way as to produce an analysis PDF with minimum variance. One can see that if  $\mathbf{G} = \underline{\mathbf{0}}$ , equation (2.26) would reduce to the expression obtained above in equation (2.16). With this more general expression for  $\mathbf{K}$ , the update equation for the analysis PDF's covariance changes slightly:

$$\mathbf{P}^a = \mathbf{P}^f - \mathbf{K} (\mathbf{P}^f \mathbf{H}^T - \mathbf{G})^T. \quad (2.27)$$

### Reusing observations in the Kalman filter

The forms of equations (2.26) and (2.27) are satisfying, for if one tries to reuse observational information, that is, to further update an analysis PDF by the same observations' PDF *taking into account their error correlations*, then one finds the Kalman filter gives zero further increment. To see this, first evaluate the expected value of  $\mathbf{G}$  when further updating an analysis by the same observations. Right multiply equation (2.24) by  $\boldsymbol{\epsilon}^o{}^T$ , and then apply the expectation operator. Doing so yields:

$$\mathbf{G}^a = \mathbf{K} \mathbf{R}, \quad (2.28)$$

where the superscript  $a$  denotes that  $\mathbf{G}$  is the correlated error between a previous analysis and the observations. To show that this gives zero increment, plug this expression into equation (2.26):

$$\begin{aligned} \mathbf{K}^a &= (\mathbf{P}^a \mathbf{H}^T - \mathbf{G}^a) (\mathbf{H} \mathbf{P}^a \mathbf{H}^T + \mathbf{R} - \mathbf{H} \mathbf{G}^a - \mathbf{G}^{aT} \mathbf{H}^T)^{-1} \\ &= ((\mathbf{I} - \mathbf{K} \mathbf{H}) \mathbf{P}^f \mathbf{H}^T - \mathbf{K} \mathbf{R}) \dots \\ &= (\mathbf{P}^f \mathbf{H}^T - \mathbf{K} (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})) \dots \\ &= (\mathbf{P}^f \mathbf{H}^T - \mathbf{P}^f \mathbf{H}^T) \dots = \underline{\mathbf{0}}. \end{aligned}$$

The machinery correctly predicts no further update as all the available information in the observations has already been used. Had one assumed  $\mathbf{G}^a = \underline{\mathbf{0}}$ , the Kalman filter would lead to an estimate biased toward the observations. We return this point in the last chapter.

## The Kalman filter forecast step

Mention should also be made of the forecast step of the Kalman filter. Just as the update step of the filtering process can be derived from Bayes rule under the assumptions of Gaussian PDFs, so can the forecast step from stochastic dynamic prediction (see section 2.3.2) under the assumption of linear dynamics. If the dynamics are linear, then we can replace the operator  $\mathcal{M}[\cdot]$  with a matrix  $\mathbf{M}$ , such that  $\mathbf{M}(\mathbf{x} + \boldsymbol{\varepsilon}) = \mathbf{M}\mathbf{x} + \mathbf{M}\boldsymbol{\varepsilon}$ . If the correct dynamics are linear, then:

$$\mathbf{x}_i^t = \mathbf{M}\mathbf{x}_{i-1}^t \quad (2.29)$$

$$\mathbf{x}_i^f = \mathbf{M}\mathbf{x}_{i-1}^a, \quad (2.30)$$

where the subscripts denote successive time steps. Equation (2.30) explicitly shows the cycling nature of filtering, where the analysis from the last time DA was performed serves as the initial state for the next forecast. Subtracting equation (2.29) from equation (2.30) yields:

$$\boldsymbol{\varepsilon}_i^f = \mathbf{M}\boldsymbol{\varepsilon}_{i-1}^a. \quad (2.31)$$

An evolution equation for the covariance matrix comes from taking the outer product of each side of the above equation with itself and applying the expectation operator. Doing so yields:

$$\mathbf{P}_i^f = \mathbf{M}\mathbf{P}_{i-1}^a\mathbf{M}^T. \quad (2.32)$$

This is a simplification from the stochastic dynamic prediction approach because linearity was assumed. As stated in section 2.3.2, the evolution equation for each moment depends on each successive higher moment, yet there is no reference to  $\mathbf{P}^a$  in equation (2.30). This is because linear dynamics removes its contribution.

For completeness, we briefly mention the traditional inclusion of a model error term. If one believes model error to be present, then an additional term should be present in equation (2.29), a vector noise process  $\boldsymbol{\eta}_{t-1}$ . If this is the case, then equation (2.31) should also include  $\boldsymbol{\eta}_{t-1}$ , and equation (2.32) should include another term,  $\mathbf{Q} \equiv \langle \boldsymbol{\eta}_{t-1}\boldsymbol{\eta}_{t-1}^T \rangle$ . In the perfect model scenario,  $\mathbf{Q} = \mathbf{0}$ .

## Extensions to nonlinearity & non-normality

There are three major linear aspects to the Kalman filter. First and foremost is the fact that the analysis PDF mean is posed as a linear combination of the means of two other PDFs. As shown above in equation (2.11), this is a natural consequence of the prior PDF

and the observations' PDF both being Gaussians, but if one were to apply the Kalman filter to problems where the PDFs were not necessarily Gaussians, then the linear combination approach would not be strictly correct. The other two aspects of linearity are associated with the two operators that have been assumed to be linear when in reality they need not be (and often are not). The observation operator,  $\mathcal{H}[\cdot]$ , is possibly linear in the case of fixed platform, direct observations made at locations in space that correspond to gridpoints in the discretized model domain, but otherwise will not be. The dynamical operator,  $\mathcal{M}[\cdot]$  (i.e., the model), is rarely going to be linear; if one is predicting a chaotic system, then it will never be linear. As stressed above, it is with frustration that the correct approach to probabilistic prediction and its direct approximations cannot be implemented in most cases; however, progress has been made in forging nonlinear extensions to these.

In regards to the observation operator and the model, careful extensions can be made via Taylor series expansions. The approach in deriving both the Kalman filter update and forecast step relied on being able to subtract  $\mathbf{x}^t$  from both sides of equations governing state vectors to yield equations governing their errors (e.g., equations (2.24) and (2.31)). The validity of this subtraction depends on the validity of  $\mathbf{H}(\mathbf{x} + \boldsymbol{\varepsilon}) = \mathbf{H}\mathbf{x} + \mathbf{H}\boldsymbol{\varepsilon}$  and  $\mathbf{M}(\mathbf{x} + \boldsymbol{\varepsilon}) = \mathbf{M}\mathbf{x} + \mathbf{M}\boldsymbol{\varepsilon}$ . When  $\mathbf{H}$  and  $\mathbf{M}$  are written as matrices, these relations are true by definition of linearity. However, one can also satisfy the relations by way of validly linearizing the operators. First one expands each operator in a Taylor series about a state vector:

$$\mathcal{H}[\mathbf{x} + \boldsymbol{\varepsilon}] = \mathcal{H}[\mathbf{x}] + \mathbf{H}\boldsymbol{\varepsilon} + \frac{1}{2}\boldsymbol{\varepsilon}^T \widehat{\mathbf{H}}\boldsymbol{\varepsilon} + \dots \quad (2.33)$$

$$\mathcal{M}[\mathbf{x} + \boldsymbol{\varepsilon}] = \mathcal{M}[\mathbf{x}] + \mathbf{M}\boldsymbol{\varepsilon} + \frac{1}{2}\boldsymbol{\varepsilon}^T \widehat{\mathbf{M}}\boldsymbol{\varepsilon} + \dots, \quad (2.34)$$

where here  $\mathcal{H}$  and  $\mathcal{M}$  are the full nonlinear operators originally introduced,  $\mathbf{H}$  and  $\mathbf{M}$  are the Jacobian matrices of the nonlinear operators evaluated about the state  $\mathbf{x}$ , and  $\widehat{\mathbf{H}}$  and  $\widehat{\mathbf{M}}$  are the Hessian tensors of the nonlinear operators evaluated about the state  $\mathbf{x}$ . A function's Jacobian is the matrix of first derivatives of the function's output vector with respect to its input vector; its Hessian is a third order tensor of the second derivatives of the function's output vector with respect to its input vector. If the Hessian terms and all higher order terms in the above Taylor series can be neglected, then the nonlinear operators can be said to be linearized. In general, discounting very pathological cases, the neglect of terms higher than the Jacobian matrices can be made valid by choosing small enough errors,  $\boldsymbol{\varepsilon}$ . In the model's case, there is a related, added constraint on the integration time being small enough (i.e., to ensure that  $\boldsymbol{\varepsilon}$  stays small enough). The model's Jacobian matrix,  $\mathbf{M}$ , is called the

*tangent linear model.* The tangent linear model gives the local tangent directions in state space to the full nonlinear state trajectory given by  $\mathcal{M}[\mathbf{x}]$ .

### The extended Kalman filter

By assuming that  $\mathcal{H}$  and  $\mathcal{M}$  can be linearized, then the above derivations can still proceed. The results change in that the full nonlinear operators are used when treating the PDFs' means (i.e., in the innovation vector and in evolving the analysis PDF's mean), but their Jacobian matrices are used when treating the PDFs' covariances. Including fully nonlinear operators disallows statements of strict optimality. The resulting equations are termed the extended Kalman filter (EKF), and for convenience, they are summarized here (under the common assumption that  $\mathbf{G} = \mathbf{0}$  and the perfect model assumption that  $\mathbf{Q} = \mathbf{0}$ ):

$$\mathbf{x}_i^f = \mathcal{M}[\mathbf{x}_{i-1}^a] \quad (2.35)$$

$$\mathbf{P}_i^f = \mathbf{M}_{i-1} \mathbf{P}_{i-1}^a \mathbf{M}_{i-1}^T \quad (2.36)$$

$$\mathbf{x}_i^a = \mathbf{x}_i^f + \mathbf{K}_i \left( \mathbf{y}_i^o - \mathcal{H}[\mathbf{x}_i^f] \right) \quad (2.37)$$

$$\mathbf{P}_i^a = (\mathbf{I} - \mathbf{K}_i \mathbf{H}_i) \mathbf{P}_i^f \quad (2.38)$$

$$\mathbf{K}_i = \mathbf{P}_i^f \mathbf{H}_i^T \left( \mathbf{H}_i \mathbf{P}_i^f \mathbf{H}_i^T + \mathbf{R}_i \right)^{-1}, \quad (2.39)$$

where the subscripted  $i$ 's on  $\mathbf{H}$  and  $\mathbf{M}$  denote that they have been linearized about the state at time  $i$ . As pointed out by Cohn (1993), the EKF is not quite consistent with a Gaussian closure assumption in stochastic dynamic prediction. As noted above, stochastic dynamic prediction has an evolution equation for the analysis PDF mean that is different than the evolution equation for a specific state of the system. The EKF is missing a term in equation (2.35) proportional to the model's Hessian tensor,  $\hat{\mathbf{M}}$ . This neglect is justified if the tangent linear approximation is truly valid, but becomes the matter of some concern if it is not. There are many other nonlinear extensions to the Kalman filter (e.g., Jazwinski 1970; Wunsch 1996), but none are as popular as the EKF. Even though the EKF does not deliver a strictly optimal estimate of the analysis PDF's mean vector (because of nonlinearity and its ability to impart non-normality), if the assumed linearizations are valid, then it should not be too different. Therefore, it is really the analysis PDF's covariance, given by equation (2.38) that is the sought statistical descriptor of the estimate.

## Ensemble-based Kalman filtering

Even though the Kalman filter and its nonlinear extensions are heavily simplified versions of the correct approach, even they are difficult to implement in large-scale prediction problems. Computation-wise, a state vector with  $n$  elements has a covariance matrix with  $n(n+1)/2$  unique elements (because covariance matrices are symmetric). When  $n \sim 10^7$ , this becomes prohibitively large even to store, let alone perform the matrix multiplication required in equation (2.36). Information-wise, it is unclear how one would deliberately and meaningfully specify all the  $10^{14}$  elements in the covariance matrix. Hence, for both of these reasons, simplifications of the full EKF machinery that approximate the analysis mean and covariance have been sought. *The myriad simplifications can be grossly grouped together as truncated filters in that they effectively operate in a dimension much smaller than  $n$ .* Specific schemes will be considered in later chapters, but as this thesis is mainly concerned with the class of simplifications based on Monte Carlo methods, their basics are discussed here.

As stated before in section 2.3.2, strict Monte Carlo methods are valid approximations of the correct problem, but are also prohibitively expensive. It was Evensen (1994) that first conceived of blending the restrictive assumptions of the EKF with Monte Carlo methods to yield an implementable DA system that was based, at least in principle, on the correct approach. In his so-called ensemble Kalman filter (EnKF), he let an ensemble of forecasts diverge naturally under the system's nonlinear dynamics so as to approximate the full nonlinear evolution of the analysis PDF, not just its first two moments. However, he only used the first two approximate moments from the ensemble to inform the Kalman filter update mechanics. That is to say, the prior PDF's mean and covariance in the filter are approximated by the ensemble population mean and covariance. The implications of updating only selected moments of a generally non-Gaussian PDF will be considered in the next chapter. In the EnKF, each ensemble member is updated individually so that their analysis mean and analysis covariance closely approximate the desired values given by equations (2.37) and (2.38).

Notation-wise, if one has an ensemble of  $N$  state vectors, denoted  $\mathbf{z}_j$ , for  $j = 1, N$ , then the unbiased population mean and covariance are, respectively:

$$\bar{\mathbf{z}} = \frac{1}{N} \sum_{j=1}^N \mathbf{z}_j \quad (2.40)$$

$$\bar{\mathbf{P}} = \frac{1}{N-1} \sum_{j=1}^N (\mathbf{z}_j - \bar{\mathbf{z}}) (\mathbf{z}_j - \bar{\mathbf{z}})^T. \quad (2.41)$$

It proves useful to consider the Reynolds decomposition of the ensemble into its mean and perturbation,  $\mathbf{z}_j = \bar{\mathbf{z}} + \mathbf{z}'_j$ .<sup>5</sup> Equation (2.41) is often rewritten after defining an ensemble matrix,  $\mathbf{Z}$ , whose columns are the ensemble vectors. The matrix-wise decomposition is then  $\mathbf{Z} = \bar{\mathbf{Z}} + \mathbf{Z}'$ , where all the columns of  $\bar{\mathbf{Z}}$  are the ensemble mean, and the columns of  $\mathbf{Z}'$  are the ensemble perturbation vectors. Equation (2.41) can then be rewritten,  $\bar{\mathbf{P}} = \frac{1}{N-1} \mathbf{Z}' \mathbf{Z}'^T$ .

In an ensemble context, the forecasting step of the EKF, equations (2.35) and (2.36), are replaced by  $N$  integrations of the system dynamics, one for each ensemble member. If nonlinear error growth becomes important (i.e., the tangent linear approximation fails), the aforementioned issue noted by Cohn (1993) of neglecting the model's Hessian is not a worry because the system dynamics is evolving actual states, not a PDF's mean. As stated above, the crux of the EnKF is to not explicitly evolve  $\mathbf{P}^f$  but rather approximate it by the dispersion of the ensemble forecasts. The update step is still based on the Kalman filter notion of forming a linear combination of information, and therefore in an ensemble context, one is free to treat the update of the ensemble mean as separate from the update of the ensemble perturbations. Following the framework of Whitaker and Hamill (2002), one can pose the EnKF update as:

$$\bar{\mathbf{z}}^a = \bar{\mathbf{z}}^f + \mathbf{K} \left( \bar{\mathbf{y}}^o - \overline{\mathcal{H}[\mathbf{z}^f]} \right) \quad (2.42)$$

$$\mathbf{z}'_j{}^a = \mathbf{z}'_j{}^f + \tilde{\mathbf{K}} \left( \mathbf{y}'_j{}^o - (\mathcal{H}[\mathbf{z}^f])'_j \right), \quad (2.43)$$

where in the innovation vectors, one really does often consider the mean vector  $\overline{\mathcal{H}[\mathbf{z}^f]}$  as opposed to  $\mathcal{H}[\bar{\mathbf{z}}^f]$ . One need not do this in principle, but doing so allows one to avoid linearizing the observation operator when approximating the gain matrix — its necessary matrix products can be estimated directly from the “observed ensemble.” For example,

$$\mathbf{H} \mathbf{P}^f \mathbf{H}^T \approx \frac{1}{N-1} \left( \left( \mathcal{H}[\mathbf{Z}^f] - \overline{\mathcal{H}[\mathbf{Z}^f]} \right) \left( \mathcal{H}[\mathbf{Z}^f] - \overline{\mathcal{H}[\mathbf{Z}^f]} \right)^T \right),$$

where the notation has been liberally generalized to show that the observation operator has been applied to each column of the ensemble matrix,  $\mathbf{Z}^f$ . Because of the over-arching linear combination of Kalman filter based methods, this approximation will be suspect whenever  $\overline{\mathcal{H}[\mathbf{z}^f]}$  is not approximately equal to  $\mathcal{H}[\bar{\mathbf{z}}^f]$ . It can be shown via a Taylor series expansion of  $\mathcal{H}[\mathbf{z}^f]$  about  $\bar{\mathbf{z}}^f$  that the equivalency of these two terms depends on having the size of the ensemble perturbations be small compared to the length scales of the ensemble mean.

---

<sup>5</sup>We use “ $\mathbf{z}$ ” to denote ensembles by tradition. In general,  $\mathbf{z}$  denotes an ensemble related quantity, whereas  $\mathbf{x}$  denotes either a deterministic quantity or a generic state vector.

The framing of ensemble filtering by equations (2.42) and (2.43) is somewhat general, as it allows for the perturbations to be updated by a different weighting matrix,  $\tilde{\mathbf{K}}$ , than the mean. Also, the convention of writing the observation vector as a mean with ensemble perturbations allows for Monte Carlo simulation of the observations' PDF. There are several ways to frame an ensemble-based Kalman filter, and the details determine the treatment of  $\tilde{\mathbf{K}}$  and  $\mathbf{y}_j^o$ . Broadly grouped, ensemble-based Kalman filters have taken either *stochastic* approaches or *deterministic* approaches. The specifics of the “state-of-the-art” formulations will be considered in depth in the next chapter.

### Covariance matrix considerations

By examining equation (2.41), it is clear that if  $N < n$ , then the matrix  $\bar{\mathbf{P}}$  can never be full rank; similarly, just by statistical averaging, if there is sizable redundancy within the ensemble, then  $\bar{\mathbf{P}}$  may not be full rank even if  $N \geq n$ . If the ensemble members were initially generated sensibly, then as  $N \rightarrow \infty$ ,  $\bar{\mathbf{P}} \rightarrow \mathbf{P}$ . However, the reason that ensemble filters enable implementation is that one attempts to succeed with  $N \ll n$ . It is reasonable to question the meaning of severely rank deficient covariance matrices and the validity of working with them. There are several responses one might give. *The first point to make is that, mechanically, there is nothing problematic with treating a singular prior covariance matrix.* In fact, as is evident in the expression for the Kalman gain matrix, equation (2.39),  $\mathbf{P}^f$  itself is not inverted. Indeed, the matrix inversion is done in observation space, and can only be singular if both  $\mathbf{H}\mathbf{P}^f\mathbf{H}^T$  and  $\mathbf{R}$  are singular, and  $\mathbf{R}$  is not likely to be singular as it corresponds to the uncertainty information of the available observations. However, if one were to begin with the least squares cost function in equation (2.12), then the inverse of  $\mathbf{P}^f$  is needed and one will not be able to proceed in a straightforward manner if it is singular.

Two common ways to make a covariance estimate full rank, should one find need to, are either to blend the flow-dependent information that  $\bar{\mathbf{P}}$  gives with a full-rank climatological  $\mathbf{P}^f$  (e.g., Hamill and Snyder 2000) or to apply a user-specified localization to the covariance matrix. Localization was originally proposed as a method to eliminate spurious long-distance correlations (e.g., a coincidental significant correlation of the temperature in Boston with the temperature in Beijing due to a relatively small ensemble), however, one of its consequences is to force a full-rank covariance matrix by imposing an external decorrelation length scale to the problem. Houtekamer and Mitchell (1998) were among the first to suggest this for ensemble-based filtering; Gaspari and Cohn (1999) devised a mathematically rigorous localization to ensure that the localized covariance matrix would be full rank and positive definite; recently, Anderson (2005) suggested a hierarchical approach borrowed

from turbulence closure studies to allow an ensemble of ensembles determine the “correct” localization structure for covariance matrices. Clearly, both hybrid blending and localization are crutches necessitated by attempting to use a small number of ensemble members: both have a degree of arbitrariness and their need should disappear as  $N \rightarrow \infty$ . However, what if  $\mathbf{P}$  itself (as opposed to just  $\overline{\mathbf{P}}$ ) is singular? what is the interpretation then?

As discussed above in section 2.3.1 in reference to the mechanics of Bayes rule, non-zero covariances between elements in a state vector implies a degree of predictive power for one element given knowledge of another. If a covariance matrix is truly singular, then a valid interpretation is that one is working with too many degrees of freedom. Any complete  $n$ -dimensional system can be trivially be made singular by adding a  $n + 1$ th “dummy” variable that gives redundant information (e.g.,  $\frac{dx_{n+1}}{dt} = \frac{dx_n}{dt}$ ). One expects that this type of redundancy will never arise from a discretized system of PDEs because the original continuous equations have an infinite number of degrees of freedom, however, this does not disallow a singular covariance matrix’s existence. In fact, the philosophical underpinnings of formulating truncated Kalman filters lie on the belief that a state estimate’s covariance matrix is truly singular or at least strongly ill-conditioned. This belief springs from evidence that the effective dimension of a system is often much less than the formal dimension of a system. As discussed throughout section 2.2, the nonlinear dynamical perspective on this matter is that the effective dimension of a system *must* be less than its formal dimension simply owing to the existence of an attractor. As reviewed in that section, dissipative systems have contracting areas in their state spaces, meaning that attractors occupy zero volume in the full  $n$ -dimensional state space. As truth is bound to its system’s attractor, one expects the relevant (state dependent) dynamics to effectively occur in a subspace of the full state space; hence, the search for truncated filters.

In the case of NWP, Farrell and Ioannou (2001) argue that the effective dimension of atmospheric model error dynamics is on the order of  $10^3$ , which is  $\ll 10^7$ . Their argument is based on calculations of the spectra of Lyapunov exponents and singular values (see section 2.2.2) from NWP models: both spectra show a similar number of “significant” contributors ( $O(10^3)$ ) compared to the formal dimension of the models ( $O(10^7)$ ). If this is true, then one may expect nearly singular covariance matrices when represented in the full formal dimension  $n$ . Consequently, one may expect to be able to get away with representing the dynamics of interest with fewer numbers of degrees of freedom (e.g.,  $N \ll n$ ).

## 2.4 Nonlinear dynamical approach to NWP

### 2.4.1 The relevance of attractors

As Bjerknes (1904) anticipated, NWP can be treated as a nonlinear dynamical system and predicted as such. However, what he did not know at the time was that this treatment is very involved. Indeed, atmospheric dynamics appears to be very sensitive to its initial conditions and may even have a strange attractor. If the atmosphere is chaotic, then it truly ought to be approached probabilistically. Are atmospheric dynamics chaotic? Before addressing that, it is even worth questioning whether atmospheric dynamics are deterministic. Whether atmospheric dynamics are truly deterministic is somewhat of a philosophical question. One inevitably butts heads with questions like whether quantum mechanics is relevant to atmospheric dynamics and at what height the atmosphere ends. *Without answering those questions here, we simply point out that to within the approximations made in the governing equations, the system is deterministic.*<sup>6</sup> However, it is unknown whether the true atmosphere has an attractor or whether any discretized numerical models of it have attractors. In his work on the matter, Lorenz has come to believe that dissipative NWP models do have attractors, but that any such representation of the real atmosphere's attractor would be "fuzzy" (personal communication, Lorenz 2004).

There have been many attempts to reconstruct attractors from atmospheric data sets and atmospheric model output, but such studies are ultimately inconclusive (e.g., Tsonis and Elsner 1989; Tiwari and Rao 1999). However, the evidence at hand is consistent with the atmosphere and its model representations being chaotic and having attractors. For instance, predictions are observed to be sensitive to their initial conditions, and the state space appears to be bounded (at least until the sun explodes). Also, models of different resolution are capable of giving broadly similar analysis results, and, as described at the end of section 2.3.3, the spectra of estimated Lyapunov exponents and singular values from NWP models are similarly distributed with a relatively few number of important growing structures. This suggests that the systems all have an effective dimension much less than their formal dimension (though this is an unclear concept for the atmosphere itself). Certainly, it does not seem likely that we shall ever know the true state of the system (or have access to the true system dynamics). Based on the evidence at hand, we conclude that we cannot reject the null hypothesis that the atmosphere and its model

---

<sup>6</sup>Some prediction efforts employ stochastic procedures like stochastic parameterizations (e.g., Buizza et al. 1999), but these are attempts to compensate for model errors, not consequences of intrinsic stochasticity within the model equations

representations are chaotic. Hence, we carry forth as if they are chaotic and seek methods for their state dependent, probabilistic state estimation and prediction.<sup>7</sup> Therefore, having a correct approach and understanding its approximations is instructive. To gauge the chasm we must ford, we next review the DA schemes being used in major operational centers.

## 2.4.2 Operational approaches

Though there is wide acknowledgment that NWP should be treated probabilistically, none of the major operational NWP centers are performing this procedure in a consistent fashion. This has much to do with the fact that operational NWP outfits have slightly different goals than scientists who study predictability — NWP operations care first and foremost about making their predictions more accurate, and this can lead to a fair bit of engineering where methods are devised to improve skill scores that have little or no theoretical motivation. Essentially, if arbitrary method X works, then arbitrary method X is often what is used. This principle is perhaps best seen in how the various models are tuned. Scientists care more for soundly improving the approaches to prediction based on lessons learned from predictability studies. This slight conflict of motivation along with a wide chasm of appropriated resources (e.g., operations having dedicated supercomputers for their models) has made it difficult for results from academic predictability studies to influence operational prediction practices.

### Ensemble-based methods

One clear advance toward probabilistic prediction has been the addition of an Ensemble Prediction System (EPS) to most of the world's major operational centers. However, to date, the ensembles have only been used for forecasting information; no one has yet to use the information in the EPS forecasts to inform the DA step. The Canadian Meteorological Centre (CMC) has plans to implement an operational EnKF (Mackenzie 2003; Houtekamer et al. 2004), but that is as close as any center has come to a full operational realization of ensemble-based DA. Interestingly, Anderson (personal communication, 2004) and Hakim (personal communication, 2004) each report obtaining results from experiments using ensemble-based DA in relatively simple models that rival the comparable analyses and forecasts from the main operational model of the United States. Anderson used a T85 version of NCAR's Community Atmospheric Model (CAM), whereas Hakim used a limited area model (LAM) over the Pacific Northwest of the United States.

---

<sup>7</sup>If one believes the atmosphere is truly stochastic, then probabilistic methods are certainly the approach to take.

Most centers have spent much time, effort, and money developing their own operational forecast models, and hence continue to produce a “control run” at a higher resolution than the members of their EPS. Since the EPS does not inform DA, the control run should be considered the mean of a presumed Gaussian prior PDF. However, no center attempts to evolve the corresponding covariance of that mean. That is to say all centers assume a static covariance matrix about their mean vector, the control run.

### **Cost function minimization (3DVar)**

The DA system used at the United States’ National Centers for Environmental Prediction (NCEP) is called the Spectral Statistical Interpolation (SSI) scheme (Parrish and Derber 1992). It essentially minimizes the cost function expressed in equation (2.12) (only with a fully nonlinear observation operator) by way of a conjugate gradient descent algorithm (a variational approach), however, it assumes that the prior PDF’s covariance matrix is static in time and that it is spectrally separable. Being static means that equation (2.36) is not used, nor is any Monte Carlo approximation thereof, to even approximately evolve  $\mathbf{P}^f$ . Being spectrally separable means that the entire covariance matrix is assumed to be a diagonal matrix when represented by spectral elements (e.g., spherical harmonics or Fourier coefficients). This ensures that  $\mathbf{P}^f$  is full-rank so that its inverse exists, but it also asserts that there is no covariance information between the various scales of the problem. One of the hallmarks of nonlinearity is that information readily flows up and down scale, and it is certain that this assumption is restrictive. Still, some simplifying assumptions must be made in order to implement an  $n$ -dimensional version of equation (2.12), and this is what the available computational resources have allowed for. The CMC also implements a version of 3DVar, but its prior PDF’s covariance information is taken from statistics compiled through the years of using its predecessor system, a so-called Optimal Interpolation (OI) scheme (essentially a local piece-wise implementation of least squares with user-specified weights). Again, the covariance matrix is static and heavily simplified.

We stress that 3DVar need not be a separate creature from ensemble-based methods. Indeed, one can try to inform the weighting matrices used in the cost function (equation (2.12)) by the state dependent information gained from an ensemble.

### **4DVar**

The European Center for Medium-Range Weather Forecasts (ECMWF) uses a similar approach to NCEP and CMC, only their initial cost function is different. The ECMWF actually performs a short term smoothing problem within the over-arching process of fil-

tering: instead of minimizing the distance between a given state to a prior estimate and newly available observations, they minimize the distance of a given state’s trajectory to a prior estimate and all the observations that have become available over the previous 12 hours. They pose the problem with their forecast model as a “hard constraint,” meaning that the minimizing trajectory must be a solution of the model. They call this 4DVar (for Four-Dimensional Variational data assimilation), but it should be clear, by extension, that this is still nothing more than the method of least squares with the added constraint that the minimizing trajectory satisfy the model dynamics (e.g., Rabier et al. 2000). Using both of the full nonlinear operators, the cost function and constraint can be written:

$$J(\mathbf{x}_0) = \frac{1}{2} (\mathbf{x}_0^f - \mathbf{x}_0)^T (\mathbf{P}_0^f)^{-1} (\mathbf{x}_0^f - \mathbf{x}_0) + \frac{1}{2} \sum_{i=0}^M (\mathbf{y}_i^o - \mathcal{H}_i[\mathbf{x}_i])^T \mathbf{R}_i^{-1} (\mathbf{y}_i^o - \mathcal{H}_i[\mathbf{x}_i]) \quad (2.44)$$

$$\mathbf{x}_i = \mathcal{M}_{0 \rightarrow i}[\mathbf{x}_0], \quad (2.45)$$

where the subscripts refer to time, with  $i = 0$  being the initial time in the 12 hour observation window considered (i.e.,  $t_0 = 12$  hours ago). The summation in the second term adds a quadratic penalty term to the cost function for each of the  $M$  batches of observations available from  $t_0$  to  $t_M$ . The model term is subscripted in equation (2.45) with time limits to show the bounds of integration. Because the dynamics are deterministic, the initially chosen  $\mathbf{x}$  determines its whole trajectory over  $[t_0, t_M]$ , therefore, as written above, one is minimizing  $J$  to find a state  $\mathbf{x}$  at  $t_0$ , but the analysis that begins the next forecast is actually that minimizing state at  $t_M$ . Because the ECMWF imposes their model as a hard constraint and because all parameters are subsumed into the operator  $\mathcal{M}[\cdot]$ , equation (2.45) can simply be substituted into equation (2.44) to yield:

$$J(\mathbf{x}_0) = \frac{1}{2} (\mathbf{x}_0^f - \mathbf{x}_0)^T (\mathbf{P}_0^f)^{-1} (\mathbf{x}_0^f - \mathbf{x}_0) + \frac{1}{2} \sum_{i=0}^M (\mathbf{y}_i^o - \mathcal{H}_i[\mathcal{M}_{0 \rightarrow i}[\mathbf{x}_0]])^T \mathbf{R}_i^{-1} (\mathbf{y}_i^o - \mathcal{H}_i[\mathcal{M}_{0 \rightarrow i}[\mathbf{x}_0]]). \quad (2.46)$$

This substitution gives an equivalent result as would have appending a *Lagrange multiplier* to the cost function, the classical approach to solving a constrained optimization problem by rendering it an unconstrained problem. Because  $\mathcal{M}$  and  $\mathcal{H}$  are nonlinear operators, minimizing equation (2.46) constitutes a nonlinear optimization problem whose solution can be very difficult to obtain. Many approximate solutions of equation (2.46) can be found

through any number of methods. One popular method is the so-called “adjoint method.” As with the cost function in equation (2.12), this cost function can be minimized by finding the trajectory that makes the gradient of equation (2.46) vanish:

$$\frac{dJ}{d\mathbf{x}_0} = 0 = -(\mathbf{P}_0^f)^{-1} (\mathbf{x}_0^f - \mathbf{x}_0) - \sum_{i=0}^M \mathbf{M}_{i \rightarrow 0}^T \mathbf{H}_i^T \mathbf{R}_i^{-1} (\mathbf{y}_i^o - \mathcal{H}_i[\mathcal{M}_{0 \rightarrow i}[\mathbf{x}_0]]), \quad (2.47)$$

where  $\mathbf{M}_{i \rightarrow 0}^T$  is the transpose of the tangent linear model, a special matrix termed the *adjoint model*. Essentially, the distance between the trajectory of  $\mathbf{x}_0$  at  $t_i$  and the observations available at  $t_i$  is computed by evolving  $\mathbf{x}_0$  forward to  $t_i$ , applying the observation operator valid at that time, and then differencing the two. However, this computed distance is valid at  $t_i$ , so to meaningfully sum its contribution to the initial state, the transposes of the linearized operators are used to transmit this information back to  $\mathbf{x}_0$ . In particular, the adjoint model linearly propagates the information back to  $t_0$ .

Though it is the most popular solution method for cost functions like equation (2.46), the adjoint method is not strictly necessary. In theory, any minimization approach can be applied (including an exhaustive search of initial condition space), but from a practical standpoint, the adjoint method seems ubiquitous for high-dimensional systems.

## Operational DA conclusion

The above is a sampling of the DA systems currently in use for NWP. If and when the CMC implements an operational EnKF, it will constitute the state-of-the-art in operational DA. It is evident that none of the systems in use today utilize state dependent error growth estimates (i.e., all centers use a static prior error covariance matrix, although some are allowed to change on a seasonal basis). As with essentially all implementable DA schemes suitable for high dimensioned systems, operational centers are forced to make error growth linearity assumptions and Gaussian error PDF assumptions.

### 2.4.3 “Physics-based” approaches

As essentially all extant DA methods suitable for NWP problems assume that error distributions are Gaussian PDFs, all can discard potentially important information contained within PDFs’ higher order moments. As described in sections 2.2.4 and 2.3.1, *this discarded information is conceptually a loss of information in state space, where DA takes place*. The work we present in chapters 4 through 6 concerns ideas of how to reconstitute some of this lost information *in physical space (i.e., the discretized model domain) based on the known system*

*dynamics*. While not necessarily posed this way, there are existing studies of other methods that have tried similar “physics-based” approaches, that is, approaches to improve DA based on knowledge of the physics of the problem at hand. Here we briefly review two of these methods: one addressing the choice of variable to assimilate and the other addresses the coordinates in which to perform the assimilation.

### **Assimilating potential vorticity**

Li et al. (1998) investigated the feasibility of essentially assimilating potential vorticity (PV) observations into a forecast model. Their study was motivated by the notion that tracer distributions are generally good proxies for PV, especially in the stratosphere, and therefore could serve as surrogate PV observations. PV would be a very nice quantity to assimilate into models because it encases the balanced dynamics of the state, a so-called “master” variable. However, as pointed out by the authors, without an externally imposed balance constraint, it is unclear whether a forecast model could ingest a new PV distribution and consistently partition it among the mass and momentum fields. Most forecast models do not use PV as a prognostic variable, but it is easily diagnosed (i.e., the observation operator is well-known). They examined how PV observations could constrain a model’s evolution. Using a shallow water model and a very simplified DA system called “nudging” (essentially a series of empirically defined corrections toward the observations), they demonstrated in a proof-of-concept fashion that such a method is possible. One could expect their results to improve further if a more realistic DA scheme were used.

### **Transforming to geostrophic coordinates**

More relevant to the problems considered within this thesis concerning coherent features, Desroziers and Lafore (1993) considered performing DA under coordinate transformations that tend to smooth out strong gradients within fluids. They began examining fronts under a geostrophic coordinate transformation. Geostrophic coordinates were generalized to three dimensions by Hoskins (1975) in his formulation of the semi-geostrophic equations. The semi-geostrophic equations are an approximation to the primitive equations that allow for frontal collapses. Geostrophic coordinates are found by transforming the standard horizontal coordinates by  $X = x + v_g/f$  and  $Y = y - u_g/f$ , where  $(u_g, v_g)$  are the components of the geostrophic wind, and  $f$  is the Coriolis parameter (assumed constant). This coordinate transform tends to render steepened wave patterns more sinusoidal, thereby effectively giving more resolution to frontal regions. Using a numerical model allowing frontal collapse and an OI DA scheme, Desroziers and Lafore were able to show their technique

gives a measurable improvement. One reason is that the smoother transformed fields more closely satisfy the DA's assumption of isotropy, and another reason is that the coordinate's dependence on the geostrophic wind gives the DA system an implicit amount of state dependence. To address more realistic problems, they had to generalize their procedure from simple geostrophic coordinates to what they call "filtered geostrophic advection coordinates," which use the real wind rather than just its geostrophic component. Constructing these special coordinates seems somewhat of an art form, but once certain user-specified parameters are chosen, it appears, at least in two-dimensional examples, that such a coordinate transform can help the state estimation process. Desroziers (1997) extended this initial study to a spherical domain with use of a least squares DA system (cost function minimization).

## 2.5 Background summary

In this background section, we have sought to reinterpret NWP and DA informed by knowledge gained from studies of nonlinear dynamical systems. The evolution of a specific state of a deterministic system can be traced as a trajectory through its state space. Prediction concerns forecasting that trajectory, and data assimilation concerns specifying the initial state. In a chaotic system monitored by noisy observations, one cannot exactly specify the initial state, and therefore the entire practice is best posed probabilistically. There are theoretically correct methods for updating and forecasting PDFs in  $n$ -dimensional state spaces, but they are not implementable in any but the most simple of systems. Consistent approximations of the correct methods exist, but even they cannot be implemented due to both computation and information constraints. Most implementable methods assume linearizable dynamics with Gaussian PDFs, somewhat restrictive assumptions for known nonlinear systems. When these assumptions are violated, the state estimation process discards important dynamical information that can lead to poor and insensible state estimates.

Ensemble-based Kalman filters seem the most promising probabilistic DA implementations to date. The community has already accepted the importance of probabilistic forecasting, as evidenced by the applications of various operational centers' EPSs. Hence, ensemble forecasts are already available. This especially makes the prospect of folding the information contained within those ensembles back into the DA process a realistic one. So, because ensemble-based DA is "state-of-the-art" and because its implementation seems imminent, we first spend time exploring the ways in which ensemble-based filters fail in the face of nonlinear error growth and non-Gaussian PDFs. The results are contained within the next

chapter.

Keeping with the historical approach of making progress in predictability within nonlinear dynamical systems via geometrical analysis, this thesis employs a hierarchy of models within which to test ideas. Low dimensioned systems allow for direct geometrical visualization, yet geophysical problems are more concerned with very high dimensioned problems. Within this thesis, low dimensioned systems are treated largely for “proof of concept,” and the high dimensioned problems are shown to be consistent with various diagnostics from the lower dimensioned systems. Because of this approach, there are multiple systems considered throughout this thesis. Instead of describing all the systems in one place, we have elected to introduce them as they arise. The systems considered are the Ikeda map, the Lorenz latitude circle model (Lorenz 95), the Korteweg-de Vries equation, a point vortex system, a barotropic basin double-gyre model, and a doubly-periodic barotropic model.



## Chapter 3

# The Impacts of Nonlinearity and Non-Gaussianity in Ensemble-Based Data Assimilation

In the previous chapter we presented a nonlinear dynamical perspective on prediction and state estimation. In particular, we stressed the utility of geometric information like a system's state space and its attractor. We impressed that the state-of-the-art implementable DA schemes are ensemble-based Kalman filters. As described in chapter 1, our interests ultimately lie in problems that obtain non-Gaussian error probability distribution functions, so we first seek to understand how the state-of-the-art methods behave as error distributions become non-Gaussian. Within this chapter we compare two ensemble-based filter formulations, a stochastic formulation and a deterministic formulation. We study and gauge the filter updates through a hierarchy of models. We find that, of course, both filter formulations can be made to fail, but that they fail in different ways. Also, we find that filters can remain accurate (by some measure) even after they have lost reliable probabilistic interpretations.

Much of this chapter is repeated verbatim from Lawson and Hansen (2004).

### 3.1 Introduction

Accurate numerical prediction of fluid flows requires accurate initial conditions. In an effort to approximate optimal estimation methods for linear problems with Gaussian error statistics, Monte Carlo methods have become a popular and realizable approach to estimat-

ing the initial conditions necessary for forecasting (e.g., Evensen and van Leeuwen 1996; Houtekamer and Mitchell 2001). Ensemble methods based on linear theory have taken two general forms: stochastic filters and deterministic filters. Both filters strive for the same result, namely to achieve the error statistics predicted by optimal linear estimation, but accomplish their goal in different fashions, the former by way of random number realizations and the latter via explicit mathematical transformations. We seek to understand the implications of these different methodologies when applied to different regimes of error growth.

Through geometric considerations in a one-dimensional example and in a two-dimensional system's state space, we formulate an interpretation of the update process for both filter types. We examine how this interpretation expresses itself in three ensemble diagnostics, root mean square (rms) analysis error, ensemble rank histograms, and measures of ensemble skewness and kurtosis. We then monitor filter behavior with these diagnostics in a hierarchy of models. By running observation system simulation experiments in a perfect model scenario with relatively large ensembles, we test filter reliability and performance in regimes with well-contained linear error dynamics and regimes with appreciable nonlinear error growth.

We find that both filters perform well, as expected, in the linear regime, but that stochastic filters are more robust *for these diagnostics* as nonlinearity becomes important. We stress that issues of robustness are highly dependent on one's choice of measure and are application dependent. For instance, the formulators of one deterministic filter algorithm, Whitaker and Hamill (2002), show that realizing the random numbers necessary for stochastic filters adds another source of sampling error to data assimilation which can easily lead to underestimation of the desired analyzed uncertainties for small to moderately sized ensembles. Similarly, Anderson (2001) points out that deterministic filters are better at retaining higher order moments through the assimilation process and allow an easier time tracking individual ensemble member trajectories. Deterministic filters are likely to be more robust if any of these features are deemed important for one's given application. This paper does not seek to proclaim that one filter formulation is better than another; it only seeks to understand the differences between them and the implications of those differences in various regimes of linearity for systems' error growth dynamics.

In section 3.2 we give a small review of the ensemble filter formulations. We begin section 3.3 by examining the filters' respective update kinematics for two hypothetical one-dimensional ensembles. We then link insight gained from this hypothetical case to three ensemble assessment methods applied to a filtering exercise in a two-dimensional chaotic sys-

tem, the Ikeda system. Then, having seen how the updates in state space express themselves in these diagnostics, we turn our attention in section 3.4 to higher dimensioned, spatially extended models, and find the same signatures in the diagnostics for those systems. Finally, we summarize our findings in section 3.5.

## 3.2 Filter formulations

No model estimate nor observation is ever perfect, and their uncertainties render data assimilation and forecasting exercises in probability. For systems following linear dynamics with errors from Gaussian probability density functions (PDFs), an optimal state estimate, in both a minimum variance sense and a maximum likelihood sense, can be achieved through use of the Kalman filter (see section 2.3.3). The Kalman filter is a recursive state estimation technique that can be derived from Bayes' rule seeking either a maximum likelihood estimate (e.g., Lorenc 1986) or a minimum variance estimate (e.g., within estimation theory, Cohn 1997), or it can be identified as a recursive least-squares problem properly weighted by the inverses of the relevant error covariance matrices (e.g., Wunsch 1996). The filter consists of a forecast step and an update step for both the mean and the covariance of a state estimate's PDF. These two moments fully characterize a Gaussian PDF.

For nonlinear systems, where PDFs can in general develop important moments not captured by a Gaussian, no such optimality conditions can be asserted. However, through the use of various linearizations, the Kalman filter framework can be retained. One such implementation is the extended Kalman filter (EKF). This filter uses the full nonlinear model to advance the state estimate (i.e., the PDF's mean), but uses a linearized version of the model to advance the state's uncertainty (i.e., the PDF's covariance). The linearized version of the model, or the tangent linear model, is re-evaluated at each time step about the latest state estimate.

Among the main challenges of implementing an EKF in geophysical problems is the computational demand. A state estimate with  $n$  elements requires a covariance matrix with  $n^2$  elements. For  $n \sim O(10^7)$ , matrix storage and arithmetic are prohibitively expensive. This computational barrier has led scientists to explore different ways to approximate the EKF's analyses. Many reduced state filters have been proposed, most relying on the notion that the actual number of dynamically important directions for error growth is much less than the full dimension of the state estimate. Examples include Verlaan and Heemink's (1997) reduced-rank square-root filter (RRSQRT), Pham et al.'s (1998) singular evolutive extended Kalman filter (SEEK), Chin et al.'s (1999) reduced order information filter (ROIF), and

Farrell and Ioannou's (2001) balanced truncation Kalman filter.

Another class of reduced state filters utilizes Monte Carlo methods. Evensen (1994) was the first to use an ensemble of forecasts to estimate the PDFs necessary in Kalman filtering, though the idea of using ensembles for probabilistic weather forecasting was posed at least twenty years earlier (e.g., Leith 1974). In Evensen's so-called ensemble Kalman filter (EnKF), the key assumption is equating the covariance of an ensemble of forecasts, each evolved with the full nonlinear model, with the linearly evolved forecast error covariance matrix necessary for the EKF. Concomitantly, one must identify the ensemble mean as one's best state estimate. This assumption allows efficient computation of the Kalman gain matrix necessary for the update step of the filters. Houtekamer and Mitchell (1998) and Burgers et al. (1998) introduced and justified the notion of using perturbed observations in an EnKF scheme. They found that in order to approximate the proper analysis error covariance, the observations must be treated as random variables. Hence, to arrive at the EnKF, the forecast step of the EKF is replaced by  $N$  ensemble forecasts, where  $N$  is the number of ensemble members, and the update step of the EKF is replaced by updating each ensemble member with the same estimated Kalman gain matrix and observations perturbed with random samples from the assumed observational uncertainty.

Equating the ensemble spread covariance with the error covariance necessary for the EKF should be a valid assumption if the ensemble is constructed in a fashion so it can sample from the important growing modes in the current state estimate. Experience has shown that the EnKF is an effective data assimilation scheme (e.g., Evensen and van Leeuwen 1996; Houtekamer and Mitchell 1998; Hansen and Smith 2000; Reichle et al. 2002). However, experience has also shown that for insufficiently large ensembles, problems may arise in implementation of these methods. Thus, compensations such as covariance localization, where one imposes a fixed correlation length scale on one's estimate of the ensemble covariance, and covariance inflation, where one artificially increases the ensemble spread about the mean, are often employed. By using relatively large ensemble sizes, we avoid implementing these practices in the work presented here.

The EnKF is not the only algorithm for a recursive ensemble filter to achieve the EKF-predicted analysis error covariance. Several independent parties have proposed equivalent variations of an ensemble filter which does not require generating perturbed observations. Because these filters do not require further random number realizations, they are termed *deterministic* filters — once the initial ensemble and the observations are known, the updated ensemble is immediately available. In contrast, the EnKF is termed a *stochastic* filter because the random numbers necessary preclude knowledge of the analyzed ensemble

until they are realized. Bishop et al.'s (2001) ensemble transform Kalman filter (ETKF), Anderson's (2001) ensemble adjustment Kalman filter (EAKF), and Whitaker and Hamill's (2002) ensemble square-root filter (EnSRF) are all deterministic filters that arrive at the desired analysis error statistics, the analysis error covariance matrix that the EKF would have retrieved, by updating the ensemble mean and then linearly transforming the ensemble members about that mean into a distribution with exactly the desired error covariance. The transformations available to the filters are rotations, translations, and rescalings in various directions.

Consider an ensemble  $\mathbf{z}_j$ , where  $j = 1, N$  and each ensemble state has dimension  $n$ . Because the analysis is formed as a linear combination of the forecast best estimate and the observations, the update step for ensemble filters can be written in terms of the ensemble mean and perturbations about that mean,  $\mathbf{z}_j = \bar{\mathbf{z}} + \mathbf{z}'_j$ . Following Whitaker and Hamill (2002), the update step for a stochastic filter, (e.g., the EnKF) can be formulated as:

$$\bar{\mathbf{z}}^a = \bar{\mathbf{z}}^f + \mathbf{K} (\mathbf{y}^o - \mathbf{H}\bar{\mathbf{z}}^f) \quad (3.1)$$

$$\mathbf{z}'_j{}^a = \mathbf{z}'_j{}^f + \mathbf{K} (\mathbf{y}'^o - \mathbf{H}\mathbf{z}'_j{}^f), \quad j = 1, N \quad (3.2)$$

$$\mathbf{P}^a = \frac{1}{N-1} (\mathbf{Z}^a - \bar{\mathbf{Z}}^a) (\mathbf{Z}^a - \bar{\mathbf{Z}}^a)^T, \quad (3.3)$$

where superscript  $a$  denotes the updated quantity, or analysis, superscript  $f$  denotes the forecasted quantity,  $\mathbf{K}$  is the approximate Kalman gain matrix, or simply a weighting matrix, found from equating the forecast ensemble covariance with the forecast error covariance,  $\mathbf{y}^o$  is an  $m$ -vector of observations,  $\mathbf{H}$  is an operator that transforms from state space to observation space,  $\mathbf{Z}$  is an  $n \times N$  matrix having the individual ensemble members as its columns,  $\bar{\mathbf{Z}}$  is a similarly dimensioned matrix with the ensemble mean in each column, and  $\mathbf{Z}'$  is a matrix whose columns contain the deviations from the ensemble mean (*cf.*, equations (2.40) and (2.41)). The vector  $\mathbf{y}'^o$  emphasizes the need for perturbed observations. The update step for a deterministic filter is similar except equation (3.2) is written

$$\mathbf{z}'_j{}^a = \mathbf{z}'_j{}^f - \tilde{\mathbf{K}}\mathbf{H}\mathbf{z}'_j{}^f, \quad j = 1, N \quad (3.4)$$

where  $\tilde{\mathbf{K}}$  is a gain matrix reserved for updating the deviations from the ensemble mean. Deterministic filters do not implement perturbed observations and, therefore, are forced to have  $\tilde{\mathbf{K}} \neq \mathbf{K}$ .

Whitaker and Hamill (2002) cite a solution for  $\tilde{\mathbf{K}}$  found in the control theory literature, and they go further to note that if the observations have uncorrelated errors, then processing

the observations serially allows one to write:

$$\tilde{\mathbf{K}} = \left( 1 + \sqrt{\frac{\mathbf{R}}{\mathbf{H}\mathbf{P}^f\mathbf{H}^T + \mathbf{R}}} \right)^{-1} \mathbf{K}, \quad (3.5)$$

where  $\mathbf{R}$  is the observational uncertainty expressed as an error covariance matrix and  $\mathbf{P}^f$  is the forecast ensemble covariance matrix. Processing observations one at a time means that the matrices  $\mathbf{R}$  and  $\mathbf{H}\mathbf{P}^f\mathbf{H}^T$  are simply scalars. This is a particularly useful form for  $\tilde{\mathbf{K}}$  as it renders their EnSRF algorithm computationally comparable to an EnKF algorithm. Bishop et al.'s (2001) and Anderson's (2001) algorithms can be shown equivalent to Whitaker and Hamill's EnSRF (Tippett et al. 2003), but they look different because those authors phrased their update processes slightly differently. Instead of solving for  $\tilde{\mathbf{K}}$ , Bishop et al. solve for  $\mathbf{T}$  such that  $\mathbf{P}^a = \mathbf{Z}'\mathbf{T}\mathbf{T}^T\mathbf{Z}'f^T$  and Anderson solves for  $\mathbf{A}^T$  such that  $\mathbf{P}^a = \mathbf{A}^T\mathbf{Z}'f\mathbf{Z}'f^T\mathbf{A}$ . *Regardless of the details, the important distinction is that these deterministic filters all achieve the proper EKF analysis error covariance statistics because a forecast ensemble is transformed to fit the expected analysis covariance exactly; stochastic filters attempt to achieve the proper analysis error covariance by blending the forecast ensemble spread with the observational uncertainty via explicitly perturbed observations.* This paper seeks to understand the implications of this key difference.

### 3.3 Geometric interpretation of ensemble filter updating

As typical geophysical problems are hugely dimensioned, it is very difficult to visualize and intuit how data assimilation methods applied to them achieve their updates. Thus, as a starting point, we examine the update process in one and two dimensions where we can easily plot state space diagrams. We first construct one-dimensional contrived ensembles to examine the filters' updates in various linearity limits. We then apply the filters to a two-dimensional dynamic system, the Ikeda map, to examine how this behavior compounds through the assimilation process. We gauge the filters via three assessment tools: rms analysis error, ensemble rank histograms, and univariate measures of ensemble skewness and kurtosis. We will show that consistent with the pictures from our kinematic example, both filters perform well in the linear regime; as nonlinearity becomes appreciable, deterministic filters break down earlier.

### 3.3.1 Update kinematics: one-dimensional example

To demonstrate how the update processes work for these two filter classes and how they behave as non-normality/nonlinearity becomes appreciable, we examine the update processes in a contrived one-dimensional example. We generate two hypothetical ensembles, one perfectly represented by a Gaussian, a linear example, and another poorly represented by a Gaussian, a nonlinear example. We then generate a synthetic observation and update the ensembles via a stochastic filter, the EnKF, and a deterministic filter, in this case the EnSRF scheme. This is the same one-dimensional Bayes rule example shown in section 2.3.1.

We first generate an ensemble with definite higher moment structure in its distribution. For this we choose a bimodal distribution comprised of two equally likely Gaussians:

$$\phi(x) = \frac{1}{2\sqrt{2\pi}} \left[ \exp\left(-\frac{1}{2}(x - \nu)^2\right) + \exp\left(-\frac{1}{2}(x + \nu)^2\right) \right]. \quad (3.6)$$

$\phi$  is the PDF and  $\nu$  is the offset of each peak from 0. For the example shown here we use  $\nu = 4$ . This PDF has zero mean and a variance of 17, which, in turn, can define a Gaussian PDF,  $N(0, 17)$ . Having defined these two PDFs, we generate random numbers from them. For the example here, we generate 5,000 ensemble members from each PDF.<sup>1</sup> Hence, we have two ensembles approximating PDFs, one Gaussian and one non-Gaussian, whose first two moments are identical. Figure 3-1 shows the two initial ensembles and the PDFs they represent. Clearly, by construction, only the top panel is sufficiently described by its first two moments.

In order to update our ensembles, observations are necessary along with an estimate of their accuracy. We suppose a hypothetical value for “truth”, 3.5, consistent with a state in the right peak of the bimodal PDF and also within one standard deviation of the Gaussian PDF’s mean. As the updated ensemble depends on the chosen observational error size, we test three cases: one where the observational error is half the ensemble’s standard deviation ( $\mathbf{R} = (\sigma/2)^2$ ), one where the errors are equal ( $\mathbf{R} = \sigma^2$ ), and one where the observational error is twice the ensemble’s standard deviation ( $\mathbf{R} = (2\sigma)^2$ ). All three cases use the same observation, which for simplicity we set equal to truth; they differ only in their confidence of this observation.

Applying an EnKF and Whitaker and Hamill’s EnSRF to both ensembles given the

---

<sup>1</sup>Clearly, no operational implementation ever approaches ensemble sizes so large, particularly in relation to the size of the “state vector”. We are not trying to mimic an operational setting here, rather we are trying to understand how these filters behave unmired by implementation constraints.

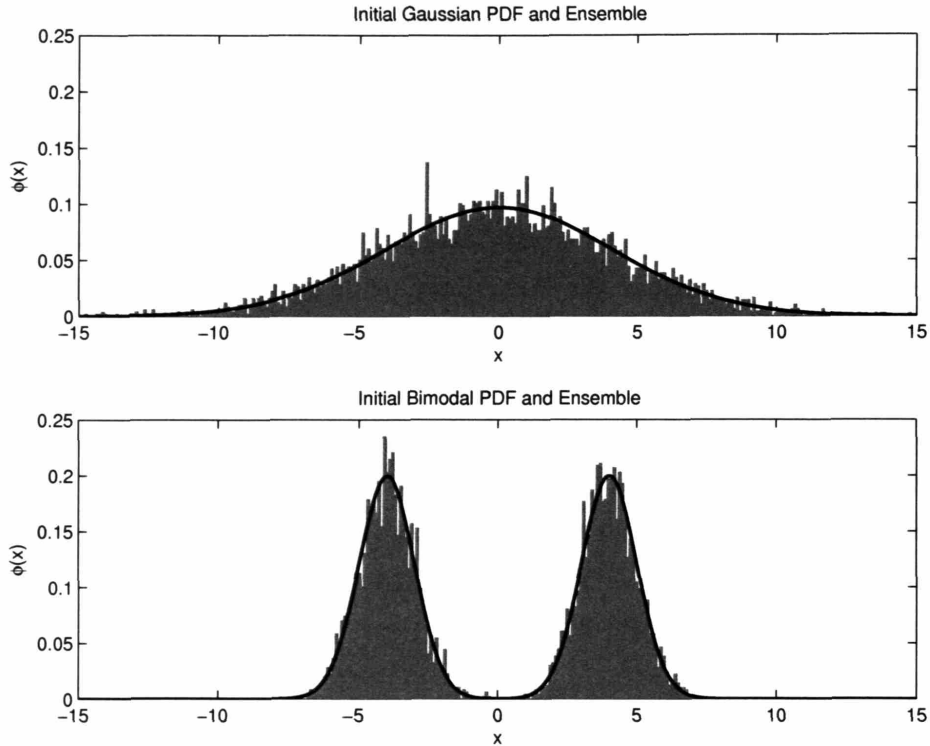


Figure 3-1: a.) The black curve is a Gaussian PDF with zero mean and variance  $\sigma^2 = 17.0$ , and the gray histogram shows an ensemble distribution of 5000 members generated from the PDF. b.) The black curve shows a bimodal PDF with the same mean and variance as the Gaussian PDF in the top panel, and the gray histogram shows a 5000 member ensemble generated from the bimodal PDF.

observation and a level of its presumed error yields a total of twelve updated ensembles. Figure 3-2 shows the six from the initially Gaussian ensemble, and figure 3-3 shows the six from the initially bimodal ensemble. The top three panels in each figure show the observation with each level of its uncertainty, and the other six panels show the updates of the initial ensembles shown in figure 3-1, one panel for each observational error/filter combination. The gray histograms show the updated ensemble distributions, and the overlain solid black lines denote the distributions predicted by application of Bayes rule to each initial PDF/observational PDF pairing (*cf.*, figure 2-2). These Bayes rule distributions are ultimately what the filters are trying to approximate.

For the initially Gaussian ensemble, the respective updated ensembles in figure 3-2 from each filter are essentially indistinguishable from one another and from the posterior distribution predicted by Bayes rule for all three observational error cases. Note the updated ensembles resemble the more accurate information source in each case (i.e., the initial ensemble or the

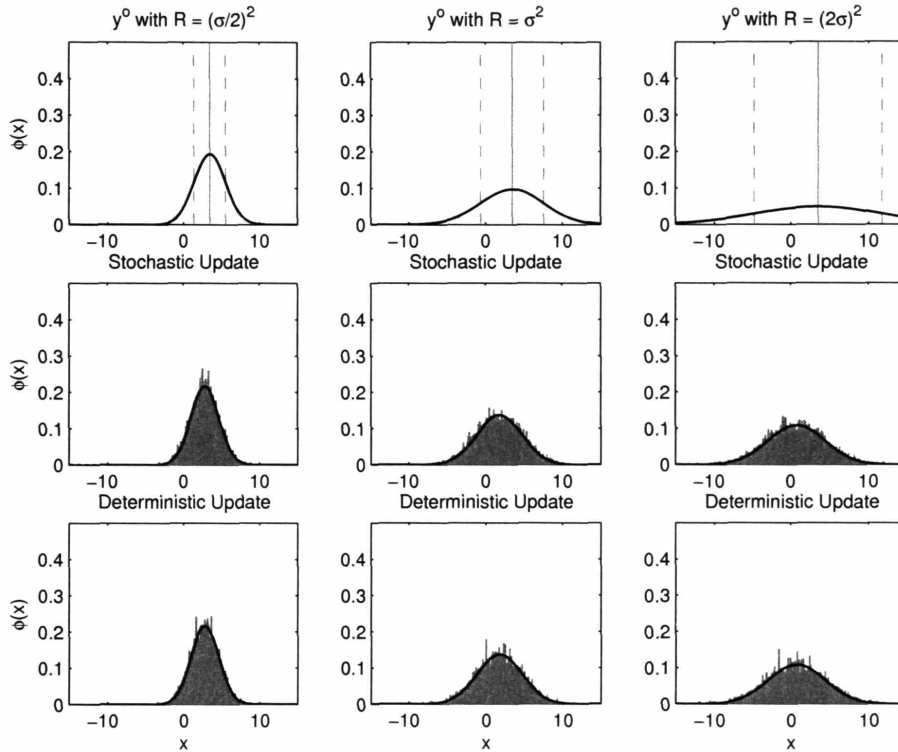


Figure 3-2: The top three panels show the observation taken with its three different accuracies, the vertical lines showing the observation  $\pm$  one observational standard deviation. The remaining six panels show the resulting ensemble updates from using either an EnKF, the top three panels, or an EnSRF, the bottom three panels, with each of the different observation accuracies. The solid black lines in the bottom six panels show the updated PDFs from application of Bayes rule.

observation). This is precisely what the update step of these filters is trying to accomplish, as the update should reflect the relative confidence in the contributing information sources.

For the initially bimodal ensemble, the respective updated ensembles in figure 3-3 are clearly different between filters, and each differ from the distributions predicted by Bayes rule. In the high accuracy observation case, the observation is accurate enough to confidently choose between the initial ensemble's two peaks. The EnKF's update has erased the bimodality, and as the observational error decreases in accuracy, more of the initial bimodality is retained. This makes sense because in the low accuracy observation case, the observational uncertainty is large enough that one cannot confidently say which of the two peaks truth is likely populating. The updated ensembles from the EnSRF, on the other hand, all show a marked bimodality, regardless of the observation's relative accuracy. This is the crux of deterministic filters, the initial ensemble is linearly transformed so that it

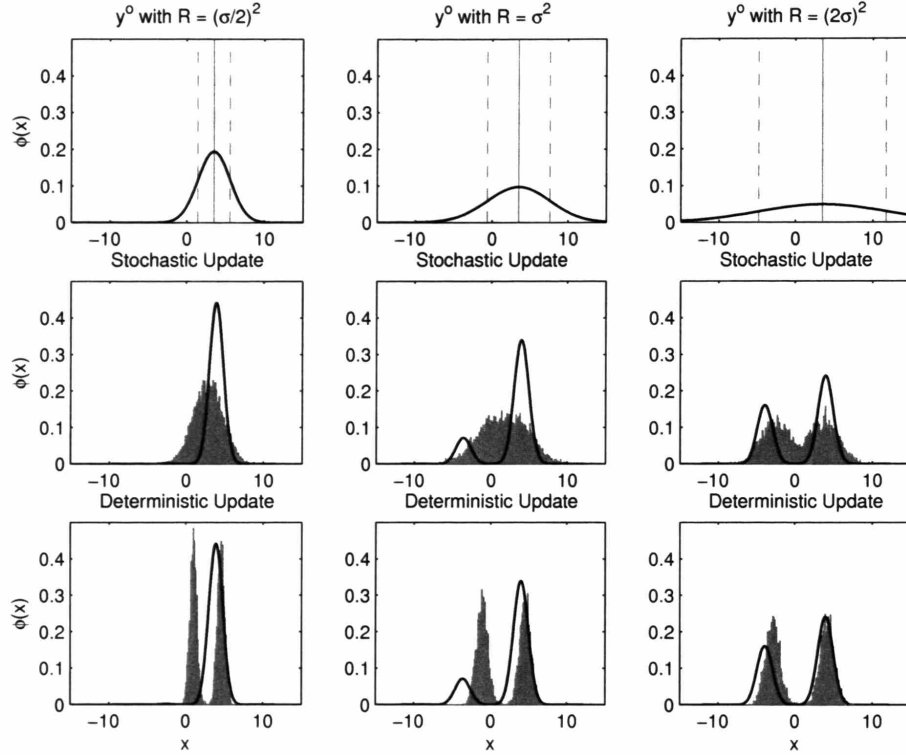


Figure 3-3: This shows the same information as figure 3-2, only for the bimodal ensemble case.

exactly fits the EKF-predicted analysis error covariance. So while the first and second moments of the distributions sampled by the updated ensemble have the desired values, they have the added feature of retaining higher moments from the initial ensemble. While this may be useful behavior in some applications, it can also be undesirable in others. For instance, note that for the high and moderate accuracy observations almost no ensemble members have the value of the observation (i.e., truth). By retaining bimodality and fitting the updated ensemble to rather accurate observations, the filter has consequently failed to sample from the most likely region for truth in the system’s assumed “state space”. Conversely, the EnSRF does not change the initial ensemble significantly when used with the low accuracy observation, and thus does not populate the initial probability gap at  $x = 0$  as the EnKF’s updated ensemble does.

These examples indicate that both filters should work very well in linear regimes. Also, inasmuch as non-Gaussian PDFs spring from nonlinear dynamics, these examples indicate that as nonlinearity becomes important, the filters can produce ensembles with members sampling regions in state space unlikely to contain truth, stochastic filters due to their

random number realizations and deterministic filters due to their retaining higher order moments. We next test these filters in a full dynamical system to confirm this suspected behavior and to gauge its consequences. We detect this behavior using ensemble assessment techniques.

### 3.3.2 Update dynamics: the Ikeda system

The Ikeda system is a two-dimensional, chaotic map originally posed as a model for the transmission of light by an optical ring cavity (Ikeda 1979). It has been used as a testbed for data assimilation applications before (e.g., Hansen and Smith 2001). The Ikeda system follows the equations:

$$x_{i+1} = 1 + \mu (x_i \cos(t) - y_i \sin(t)) \quad (3.7)$$

$$y_{i+1} = \mu (x_i \sin(t) + y_i \cos(t)) \quad (3.8)$$

$$t = a - \frac{b}{x_i^2 + y_i^2 + 1}, \quad (3.9)$$

where  $[x \ y]^T$  is the state vector, and  $a$ ,  $b$ , and  $\mu$  are parameters. The system behaves chaotically in certain parameter ranges. Figure 3-4 shows the attractor for the Ikeda system with  $a = 0.4$ ,  $b = 6.0$ , and  $\mu = 0.83$ , compiled from collecting many iterated states (the attractor in figure 2-1 uses  $\mu = 0.9$ ).

Because the Ikeda system is chaotic, states initially close together in state space will tend to diverge from one another. Hence, a good, yet imperfect estimate of truth will have a limited range of useful prediction since the initial error will grow. For small enough initial errors, the early stages of this growth generally follow linear dynamics, while the later stages begin to develop nonlinearly. Growth usually continues until the errors saturate at some climatological level, a magnitude typical of the attractor size in state space. Examining the states that comprise figure 3-4, we find a climatological mean and standard deviation of  $(\bar{x}, \bar{y}) \pm (\sigma_x, \sigma_y) = (0.66, -0.28) \pm (0.42, 0.59)$ , from which we estimate a typical climatological error size of 0.5. Note that the attractor is not Gaussian in appearance and that the climatological mean is not itself on the attractor.

In order to compare our ensemble data assimilation systems, we run observation system simulation experiments (OSSEs) in a perfect model scenario, where one freely evolving state called “truth” is used to provide synthetic observations so that an ensemble of states slightly differing from truth remains a reasonable estimate of truth’s PDF via data assimilation. This is a commonly used experimental set-up to test new DA systems and targeted




Figure 3-4: The attractor for the Ikeda system for the parameters  $a = 0.4$ ,  $b = 6.0$ , and  $\mu = 0.83$ . There are  $3 \times 2^{14}$  points plotted from many iterations.

observation strategies. To avoid implementation concerns such as covariance boosting and covariance localization, we run our OSSEs with a very large ensemble. After a spin-up period to allow the ensemble to equilibrate to the specific observing system, we employ ensemble assessment and verification techniques, all of which entail comparison to truth since it is available. Because we are comparing two data assimilation systems, our OSSEs have two ensemble estimates running in parallel. The estimates from both filters are updated using the same observations at each assimilation time.

In order to compare the filters' behavior in different linearity regimes, we need to control the degree of linearity in the system's error growth, the dynamics of which are known to be state dependent. We have two main controls available, the initial size of errors and the length of the integration over which the errors evolve. Thus, experiments considered here will differ in both observation time (i.e., the amount of time between assimilations) and observation system design. We choose to observe both state elements directly ( $\mathbf{H} = \mathbf{I}$ ) at each assimilation time; this essentially reduces varying observation system design to varying the relative accuracy of observations taken.

Since the Ikeda system is a map, the effective time step is fixed as a single map itera-

tion. In our chosen parameter range, calculations show a Lyapunov number of about 1.6, meaning one time step is about four fifths the average error doubling time, a relatively large growth per time step. Hence, the observation time is generally kept at a small number of map iterations, between 1 and 10 steps inclusive. We choose to always observe both state components with equal and uncorrelated accuracy (i.e., observational uncertainty isopleths are circles). We vary the level of observational accuracy over several orders of magnitude, from  $1.375 \times 10^{-6}$  to  $1.375 \times 10^{-2}$ . In the results shown below, all OSSEs have been performed with an ensemble size of 511. Other ensemble sizes were tested, and the results do not change significantly for larger sizes.<sup>2</sup> Verifications and statistics were compiled over 10,000 successive assimilations for each of several different initial ensemble realizations.

### Ensemble snapshots

As the Ikeda system has a two-dimensional state space, it can be easily visualized. Examining individual cases of initial ensembles versus updated ensembles reveals several general aspects of the filter updating process. The first is that ensembles are observed to disperse mainly along attractor branches. Dispersion along a primarily straight line direction in state space leads to covariance isopleths that have high eccentricity (i.e., covariance matrices have high condition numbers). Covariance isopleths have less eccentricity if there is cross-attractor branch scatter or if there is curvature in the underlying attractor. The former case is typically a product of DA pulling ensemble states off the attractor with the observation time not being long enough for all states to find the attractor before undergoing DA again. The latter case occurs whenever ensemble dispersion is comparable to the curvature length scale of the local attractor structure, for example, when the ensemble has relatively large dispersion or when the attractor has relatively sharp curvature. There are a few areas of sharp curvature readily apparent in figure 3-4, and because the attractor is fractal, we expect there to be similarly proportioned areas with large local curvature at all scales. Though these areas are relatively sparse, truth and its ensemble estimates should still visit them intermittently.

Ensemble dispersion along a straight line in state space is well approximated by a Gaussian PDF. Ensemble dispersion with cross-attractor branch scatter can be well approximated by a Gaussian PDF, especially if the scatter is symmetric about the attractor branch. Ensemble dispersion along a curved branch of the attractor will in general not be well approximated by a Gaussian PDF. Large deviations from a Gaussian PDF vio-

---

<sup>2</sup>Again, the very large ensemble size is for demonstration purposes only, not an attempt to mimic an operational setting.

late the filters' assumptions, and we thus begin to suspect problems in these situations. Common measures of departure from a Gaussian distribution are skewness and kurtosis. Skewness is a non-dimensional measure of a distribution's asymmetry, and kurtosis is a non-dimensional measure of a distribution's peakedness (i.e., how long its tails are compared to a Gaussian's). There are several common definitions for these quantities, but all entail normalizing some measure of a distribution's higher moments by an appropriate power of the distribution's standard deviation. See appendix A for details on the definitions used here. While multivariate extensions to these ideas exist (Mardia 1970), they can be quite expensive to calculate for large dimensioned systems being run with large ensemble sizes. Also, these multivariate measures are difficult to interpret and to verify. Hence, we find it more illuminating to consider the skewness and kurtosis for a single state element within the ensemble. Such measures can easily be assessed within any model for most reasonably sized ensembles. Note that having the skewness and kurtosis of a single element be consistent with a Gaussian distribution is not sufficient to assert the ensemble distribution as a whole is Gaussian, but that a single element not being consistent with a Gaussian distribution necessarily means the ensemble distribution as a whole is not Gaussian.

By calculating the skewness and kurtosis for  $x$  over our ensembles, we find that the stochastic filter (EnKF) keeps its ensemble estimates nearly Gaussian for a wide range of parameter space, whereas the deterministic filter (EnSRF) maintains estimates well-approximated by Gaussians only when the observation time is one step and the observational error is  $1.375 \times 10^{-5}$  or less. When this is the case, the observations are accurate enough to keep the estimate very near the attractor (no cross-attractor branch scatter) and the ensemble dispersion is so small that it hardly ever finds itself in a region with attractor curvature at a comparable scale. A comparison of the EnKF update and EnSRF update for such a case, with an observation time of one step and an observational error of  $1.375 \times 10^{-6}$ , is shown in figure 3-5. The top panel shows the initial ensemble (blue dots) and the updated ensemble (red dots) for the EnKF, and the bottom panel shows the same information for the EnSRF. The position of truth and the observations used to update the ensembles have been omitted, but locations on the underlying attractor are plotted as small black dots. Note the ensemble spread is contained within a state space area of  $1 \times 10^{-5}$  by  $1 \times 10^{-5}$ . Though the Ikeda system has a two-dimensional state space, the ensemble spread and updates are essentially occurring along lines.

As we allow the observation time to lengthen and the observational error to become less accurate, we find the deterministic filter's ensemble estimates readily diverge from Gaussian PDFs. Figure 3-6 succinctly summarizes the absolute value of ensemble skewness over many

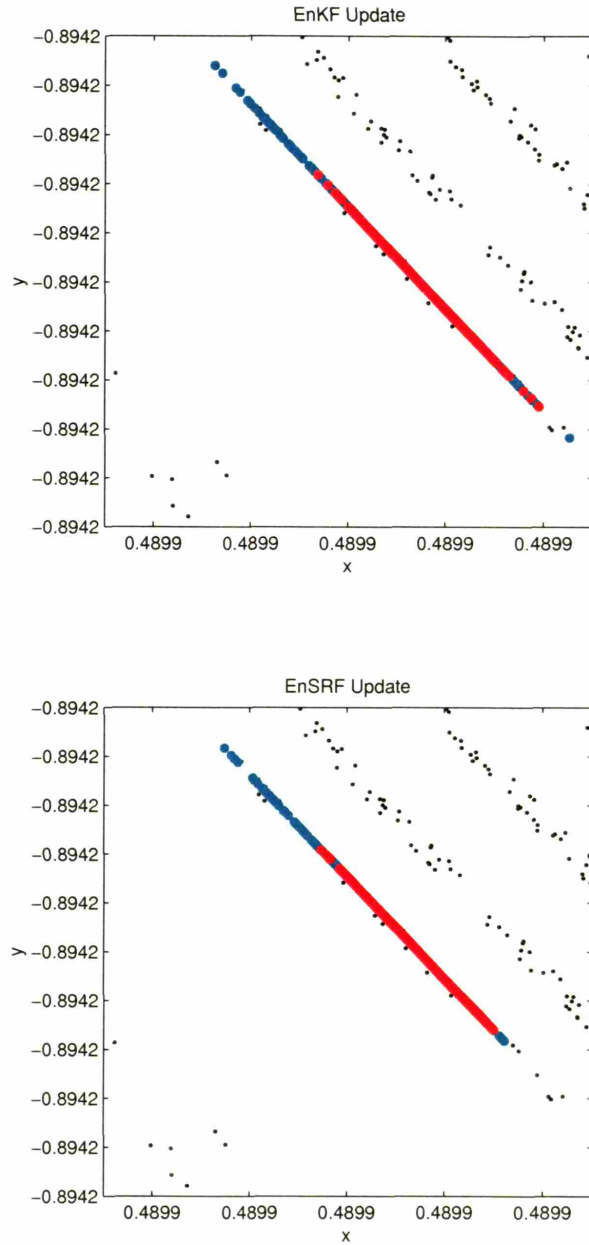


Figure 3-5: Ikeda system in a Linear Regime: a.) a snapshot from an OSSE using the EnKF in a linear regime (observation time is 1 step and observational error is  $1.375 \times 10^{-6}$ ). The small black dots are locations on the attractor, the blue dots the initial ensemble, and the red dots are the updated ensemble. b.) shows the same information, only for the EnSRF's update.

different OSSEs. Because skewness is related to an odd central moment, it can easily be positive or negative depending on what side of the mean the farthest outliers are, thereby making the mean skewness over many realizations nearly zero. Hence, we consider the

absolute value of the skewness. To characterize the values and overall distributions that ensemble skewness attains within each OSSEs time series, we employ box plots demarcating the maximum and minimum values, the median value, and the interquartile range. There are two box plots in each of the twelve figure panels, the left one is for the EnKF and the right for the EnSRF. Each panel shows the results from a different combination of observation time and observational error. The dotted line at the top of each panel shows the maximum possible value ensemble skewness can attain for the given ensemble size (for a discussion of how these maximum values are calculated, see appendix A). We omit a similar figure for kurtosis measures because they are largely redundant (i.e., the sources of large skewness are also sources of large kurtosis).




Figure 3-6: Ikeda system skewness: box plots summarizing the time series for the absolute value of ensemble skewness for each filter type over twelve different OSSEs. Each panel corresponds to a given pairing of observation time (one value for each row) and observational error (one value for each column). The left-hand box plot in each panel refers to the EnKF and the right-hand is for the EnSRF. The center line of each box plot shows the median value of the absolute value of ensemble skewness over the 10,000 assimilations and is further marked with a circle. The box's bounds show the interquartile range, and the whiskers outside the bound extend to the minimum and maximum values found over the assimilations. The horizontal dotted lines show the maximum possible value of ensemble skewness when using 511 ensemble members.

Aside from showing the departure from Gaussianity, these box plots show that in appreciable nonlinear regimes the deterministic filter's ensembles achieve skewness and kurtosis near the maximum values that ensembles their size can attain, sometimes quite frequently (e.g., the case with an observation time of three steps and an observational uncertainty of  $1.375 \times 10^{-4}$ ). This means that the ensembles are approaching or attaining a configuration where nearly all members agree on a specific value for  $x$  with just a few outliers disagreeing. Examining individual cases confirms this, as seen in figure 3-7, a snapshot from an OSSE with an observation time of 3 steps and an observational uncertainty of  $1.375 \times 10^{-3}$ . The bottom panel shows the initial and updated ensembles of the deterministic filter. The far left blue dot and red dot each constitute 510 ensemble members (all but one) occupying such close proximity that they each appear as one dot. The sole far outliers allow the ensembles as wholes to maintain fairly accurate means and to match the EKF-predicted analysis error covariance.

Clearly, the deterministic filter's ensembles in the bottom panel of figure 3-7 have highly non-Gaussian structures, and monitoring the filter's behavior over time shows that the ensembles achieve these configurations often. But the ensembles are not always in such configurations. Indeed, they do sporadically approach Gaussian PDFs as well as many intermediate configurations. One telling example is shown in figure 3-8, a snapshot from an OSSE with an observation time of 3 steps and an observational error of  $1.375 \times 10^{-2}$ . Note the initial ensembles in the two panels are similarly dispersed but that the updated ensembles are quite different. This is very similar to the one-dimensional example above where the stochastic filter tended to fill out the desired analysis covariance structure and the deterministic filter fit the ensemble exactly to that structure. In this case we see many ensemble members have been transformed well off the attractor, most especially the cluster of members already at the far right edge of the deterministic filter's ensemble. Repeated updates such as this by the EnSRF lead quickly to small groups of far outliers, which when the ensemble is forced to have the desired error covariance, necessitates that the remaining members collapse onto one another so as to counter-balance the outliers.

### **Ensemble assessment**

The three main tools utilized and presented here for assessing ensemble filter behavior are rms analysis errors, ensemble rank histograms, and univariate measures of ensemble skewness and kurtosis. Rms errors measure the accuracy of the ensemble mean as an estimate of truth. Ensemble rank histograms assess the reliability of the probabilistic estimation. And as stated above, skewness and kurtosis are standard measures for detecting departures

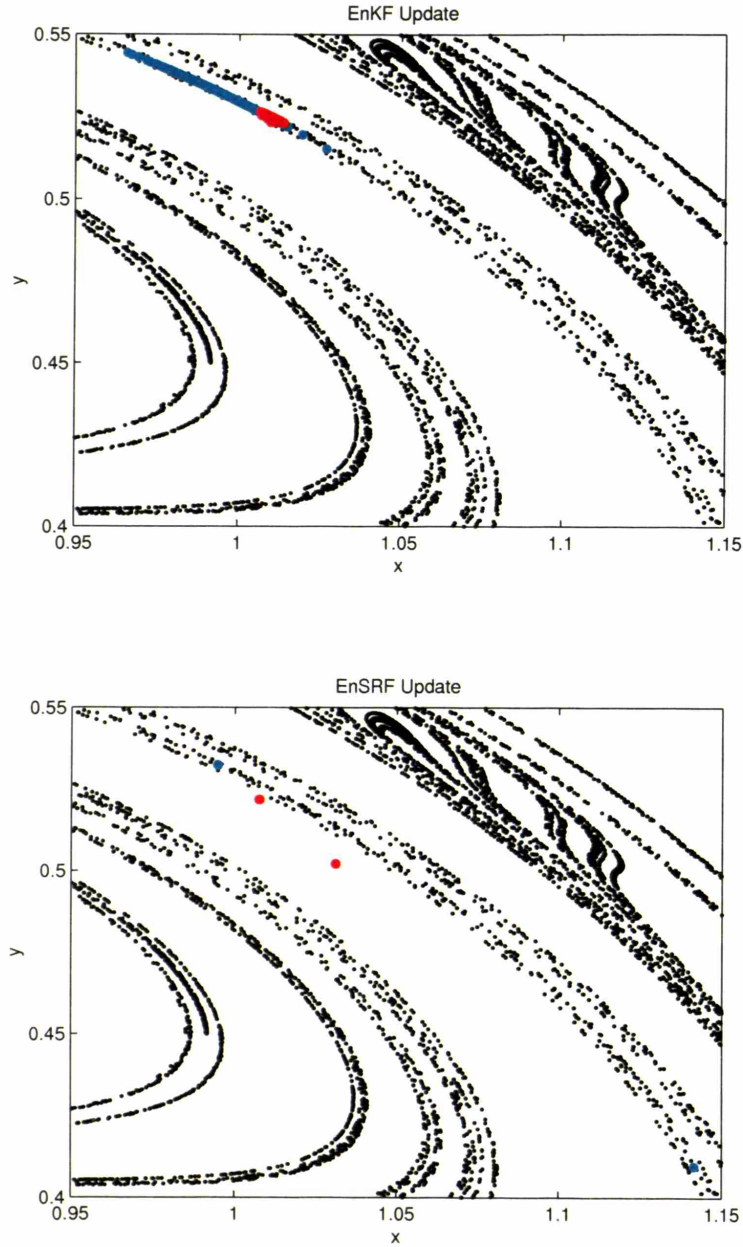


Figure 3-7: Ikeda system in a Nonlinear Regime: a.) a snapshot from an OSSE using the EnKF in a nonlinear regime (observation time is 3 steps and observational error is  $1.375 \times 10^{-3}$ ). The dots have the same meaning as in figure 3-5. b.) shows the same information, only for the EnSRF's update. The far left ensemble dots actually each contain 510 ensemble members (all but one).

from normality.

Rank histograms, also known as Talagrand diagrams, are used to assess whether truth and the ensemble all share a common PDF, that is, whether truth is statistically indistin-

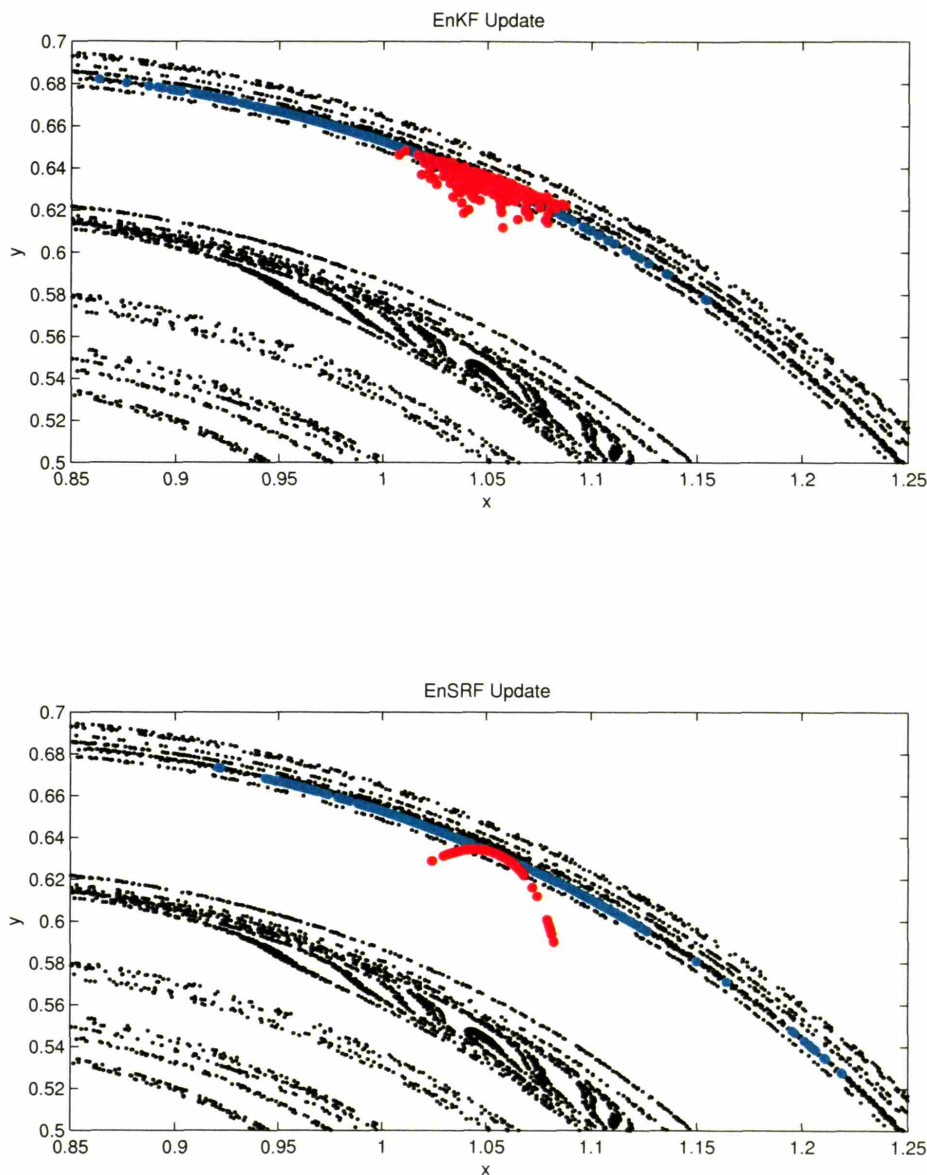


Figure 3-8: Ikeda system in a Nonlinear Regime: a.) a snapshot from an OSSE using the EnKF in a nonlinear regime (observation time is 3 steps and observational error is  $1.375 \times 10^{-2}$ ). The dots have the same meaning as in figure 3-5. b.) shows the same information, only for the EnSRF's update.

guishable from our spanning ensemble estimate. If an ensemble filter is performing properly, its members ought to be unbiased and be sampling from a probability distribution of which truth is a member. If this is the case, then probabilistic analyses and forecasts are considered reliable, for example, 70% likely events will verify 70% of the times predicted. To construct a rank histogram from an ensemble, one first chooses an index over which to sort

the ensemble, typically a single element of the state vector. Over many successive assimilations, one then finds where the verification ranks in that sorted list, and keeps tally in the bins corresponding to those ranks. If there are  $N$  ensemble members, there are  $N + 1$  bins to consider including the two end bins where verification is either greater than or less than the entire ensemble. Under the premise that the ensemble and truth are all members of the same PDF, then there should be no discernible preference to their respective ranks. Hence, reliable probabilistic data assimilation and forecasting systems should yield uniformly distributed rank histograms. Since each of the  $N + 1$  bins should be equally probable, then without loss of generality, one is free to group neighboring bins (with the caveat that each grouping must be comprised of an equal number of bins) or to sub-sample from the ensemble (e.g., only rank a fraction of the full  $N$  members). Non-uniformity in rank histograms bespeaks a number of potential problems, usually concerned with ensemble biases and spread (e.g., Hamill 2001). A multi-dimensional extension to ensemble rank histograms exists, namely minimal spanning trees (e.g., Smith and Hansen 2003), but we find these univariate measures are adequate for demonstrating the necessary points. It should also be noted that in their respective papers, Anderson (2001) and Whitaker and Hamill (2002) both chose to examine ensemble reliability by comparing the rms error of the ensemble mean to the mean of the rms errors for each ensemble member. Anderson cites the expected value this ratio ought to have if the ensemble variance is accurately representing the error in the ensemble mean. We have evaluated this ratio for several of our OSSEs and find that it confirms the behavior we observe with our assessment techniques, but being only a number, does not lend the visual interpretation that the rank histograms give.

Figure 3-9 shows the rank histograms for both the EnKF and the EnSRF applied to both a linear/Gaussian case, the left two panels, and a “slightly nonlinear”/non-Gaussian case, the right two panels. The linear case is for that shown in figure 3-5, where the observation time is 1 step and the observational uncertainty is  $1.375 \times 10^{-6}$ ; the slightly nonlinear case uses an observation time of 1 step and an observational uncertainty of  $1.375 \times 10^{-4}$ . The results from all 10,000 assimilations are displayed. In each case we are comparing the true value of  $x$  to the analyzed ensemble’s values, not to a forecast. Since we are using an ensemble size of 511, there are 512 bins to tally for the rank histograms, though we group neighboring bins such that the panels in figure 3-9 only display 32 bins. The tallies have been converted to probabilities of ensemble rank. The top two panels show the rank histograms for the EnKF. The solid horizontal line shows the expected mean value for each bin were the distribution actually uniform, and the dashed lines show the expected standard deviation from uniformity for the bin totals given that it is comprised of a fixed number of

random samples.

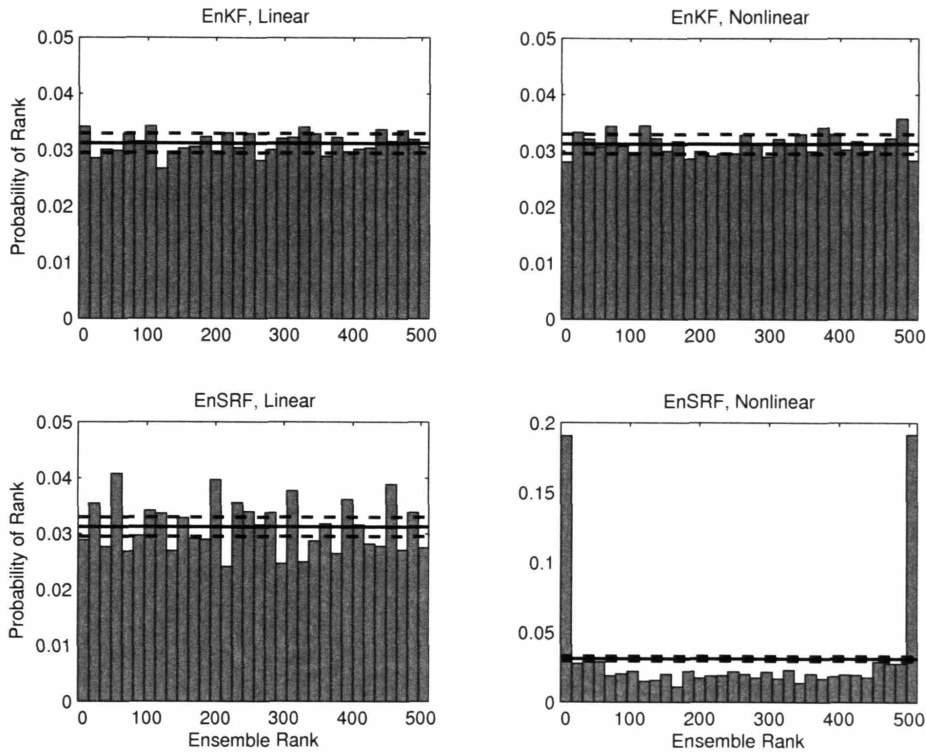


Figure 3-9: Ikeda Ensemble Rank Histograms: rank histograms are based on an ensemble size of 511, and hence each of the 32 bins shown are groupings of 16 underlying neighboring bins. Histograms are comprised of 10,000 members, and the bin counts have been converted to probabilities of rank. The solid horizontal lines show the expected value for each bin were the histograms actually uniform, and the dashed lines show the expected standard deviation from uniformity for the bin totals given that the histograms are comprised of a fixed number of random samples. The left panels show the results from a linear OSSE (observation time is 1 step and observational error is  $1.375 \times 10^{-6}$ ) and the right panels show the results from a slightly nonlinear OSSE (observation time is 1 step and observational error is  $1.375 \times 10^{-4}$ ). The top panels show the results from the EnKF and the bottom panels from the EnSRF. Note that the standard deviation lines are not strictly applicable for the EnSRF results because the tallies are not independent ranks.

Considering the linear case, the EnKF's rank histogram seems consistent with a uniform distribution as approximately two thirds of the bin totals are within the expected standard deviation. The bottom left panel shows the rank histogram for the EnSRF. Its bin totals show a larger variance in their totals. This is expected because, as cited in Anderson (2001), the tallies comprising the histogram are not independent due to the deterministic nature of the update imparting a rank correlation from one assimilation to the next. Therefore, the standard deviation lines plotted are incorrect and should reflect fewer independent degrees

of freedom. Still, there are no discernible features to make us think the ensemble is biased or not capturing the proper variability. We conclude that it also seems consistent with a uniform distribution. Inspecting the time series for rms analysis errors (not shown), we see that both filters give essentially indistinguishable analyses for all 10,000 assimilations. Similarly, as indicated by the ensemble skewness box plots in the upper left-hand panel of figure 3-6, the ensemble estimate distributions, at least for state element  $x$ , are never far from Gaussian. *From these ensemble assessment measures we conclude that in this linear range, both ensemble filters are performing accurately and reliably.*

The right two panels of figure 3-9 show the rank histograms for the slightly nonlinear case. The EnKF rank histogram is still approximately uniform, but the EnSRF rank histogram is not. The end bins of the EnSRF rank histogram are heavily populated, meaning that truth is often ranked near the ensemble edge members or outside the ensemble completely. On its own, this result is difficult to interpret beyond asserting that it is not uniform and, therefore, not sampling from the same PDF that truth is following. Hamill (2001) notes that “U-shaped” rank histograms are typically a sign that the ensemble is under-dispersive; however, we know here that the ensemble is forced to have the desired analysis error covariance, so that is not the cause *per se*. A U-shaped rank histogram is consistent with the idea of periodically spawning far outliers. This is simply because a structured, non-Gaussian distribution can generally have far outliers from its mean, and hence, even if the ensemble mean is a very good estimate of truth, many of the ensemble members may systematically be to one given side of the verification. The box plots for this case, shown in the second panel of figure 3-6, indicate that the ensembles in this parameter range are often non-Gaussian, but the measures never approach their maximum values, hence we conclude that ensemble clumping and outlier generation are probably not so severe.

*That an OSSE can give a U-shaped ensemble rank histogram yet give analyses with the expected rms error statistics indicates it is possible to form unreliable analyses that are still accurate.* This behavior does not pervade much of parameter space however: not properly sampling from truth’s PDF will eventually degrade an estimate’s accuracy. For instance, the nonlinear case shown in figure 3-8 has very similar ensemble rank histograms (not shown) to the ones for the slightly nonlinear case shown in the right-hand panels of figure 3-9. However, the time series for rms analysis error, skewness, and kurtosis are quite different between these two OSSEs, with the fully nonlinear case averaging rms errors above the observational uncertainty level and skewness and kurtosis often approaching their maximum values (see figure 3-6). These indicate that even modest departures from normality can lead to U-shaped ensemble rank histograms. Conversely, it seems a filter can maintain

appropriate levels of accuracy even when its ensembles have modestly departed from Gaussian distributions. By comparing the skewness and kurtosis measures from widely differing OSSEs using the EnKF to those using the EnSRF, we see the EnKF is able to keep its analyzed ensembles fairly close to Gaussian PDFs. Ensemble estimates from the EnSRF readily depart from Gaussian PDFs, and when they do their deviations from normality are kept through the filter’s transformations and can easily derail our probabilistic interpretation of the ensemble via ensemble rank histograms. The stochastic filter is able to produce an expected uniform rank histogram because its perturbed observations are acting sufficiently well as a Gaussian PDF repopulation device.

*By plotting full state space representations of our ensembles, we have gleaned a geometric interpretation of the ensemble filter update mechanisms. By running OSSEs with the Ikeda system, we were able to collect rms error statistics, compile rank histograms, and examine measures of skewness and kurtosis. From these we have argued that the geometric interpretation is consistent with these ensemble assessments.* Now we turn to higher dimensioned, spatially extended systems to see if this ensemble assessment behavior carries over when the filters are used in various linearity regimes.

### **3.4 Ensemble filter behavior in higher dimensioned systems**

We would like to know whether the insight gained from our simple two-dimensional experiments extends to higher dimensioned systems. Given that we cannot plot the state space for systems much over two dimensions, we look to show consistency with this insight by way of the ensemble assessment signatures discussed for the Ikeda system. We examine rms errors and ensemble rank histograms, again bolstered by skewness and kurtosis measures, for two models with differing complexity: Lorenz’s 1-latitude circle model and a two-dimensional barotropic form of the quasigeostrophic Harvard Open Ocean Model. By judiciously selecting the relevant parameters, we can place the filtering process in any desired regime along a linearity continuum. The results below generally compare linear cases to weakly nonlinear cases as they seem most relevant to real world applications, though we have explored some strongly nonlinear cases as well. The degree of nonlinear error growth can be sensibly compared between different systems, including the real world, by reference to typical error magnitudes and typical error doubling times. Ehrendorfer (1997), and references within, estimate the atmosphere’s error doubling time as  $2 \pm 0.5$  days. While we do not make strict numerical comparisons here, the success of running data assimilation cycles every 6 to 12

hours for NWP models with the given levels of observational uncertainty present suggests a similar range in which to run the experiments here, a range we may term weakly nonlinear in error growth.

### 3.4.1 Lorenz latitude circle model

The first spatially extended model we consider is Lorenz’s 1-latitude circle model, often referred to as Lorenz 95 (or sometimes 96, Lorenz 1995). It is a highly simplified model that governs the evolution of a variable around a latitude circle or a re-entrant channel. This model has been used by Lorenz and Emanuel (1998) and Hansen and Smith (2000) to explore strategies in targeted observations, and more recently by Anderson (2001) and Whitaker and Hamill (2002) as testbeds for their respective deterministic filter schemes. The model is well-documented in these references, so we refrain from repeating many details or showing a figure of a typical state, though minimum detail is included in appendix C.

For our OSSEs we use an ensemble size of  $N = 511$ , and, as with the Ikeda system, we choose to observe all state elements at each observation time. We again choose to control the degree of linearity in the problem by varying both the observation time length and the observational error size. We have performed experiments with observation times ranging from 6 hours to 48 hours (using the conversion from model time to “atmosphere time” suggested by Lorenz and Emanuel (1998) of 0.05 model time = 6 hours) and observational error scales ranging from  $2.0 \times 10^{-4}$  to 2.0. As with the Ikeda system, we perform 10,000 assimilations per OSSE. We compile rank histograms and skewness and kurtosis data based on the analyzed values for state element  $X_1$ .

Figure 3-10 shows the same rank histogram as figure 3-9, only for the Lorenz 95 system. Here the linear/Gaussian case has an observation time of 12 hours and an observational uncertainty of  $2.0 \times 10^{-4}$ . As with the Ikeda system, both filters give essentially uniformly distributed rank histograms. The skewness and kurtosis measures (not shown) confirm that neither filter is producing an ensemble that differs significantly from a Gaussian distribution. Further examination shows that both filters produce essentially indistinguishable analysis rms error time series (not shown). *Based on these measures, both filters seem to be giving accurate and reliable probabilistic analyzes.* The nonlinear/non-Gaussian case in figure 3-10 also has an observation time of 12 hours but now the observational error is 0.2. We see the deterministic filter’s rank histogram shows a similar U-shaped structure to the one we saw for the nonlinear case in the Ikeda system (lower right-hand panel of figure 3-9). Note that bins 2 and 31 are the most populated bins, as opposed to the edge bins for Ikeda. This is not too surprising as it indicates that our ensemble is still generally bounding truth, just

preferentially within either end. It stands to reason that our ensembles must be bounding truth a majority of the time as the rms analysis errors hardly differ from those of the EnKF's estimates (not shown). *However, seeing that these ensembles attain rather high values of skewness and kurtosis (not shown), then by our geometric reasoning above based on outliers it makes sense that truth often falls near the ensemble ends.*

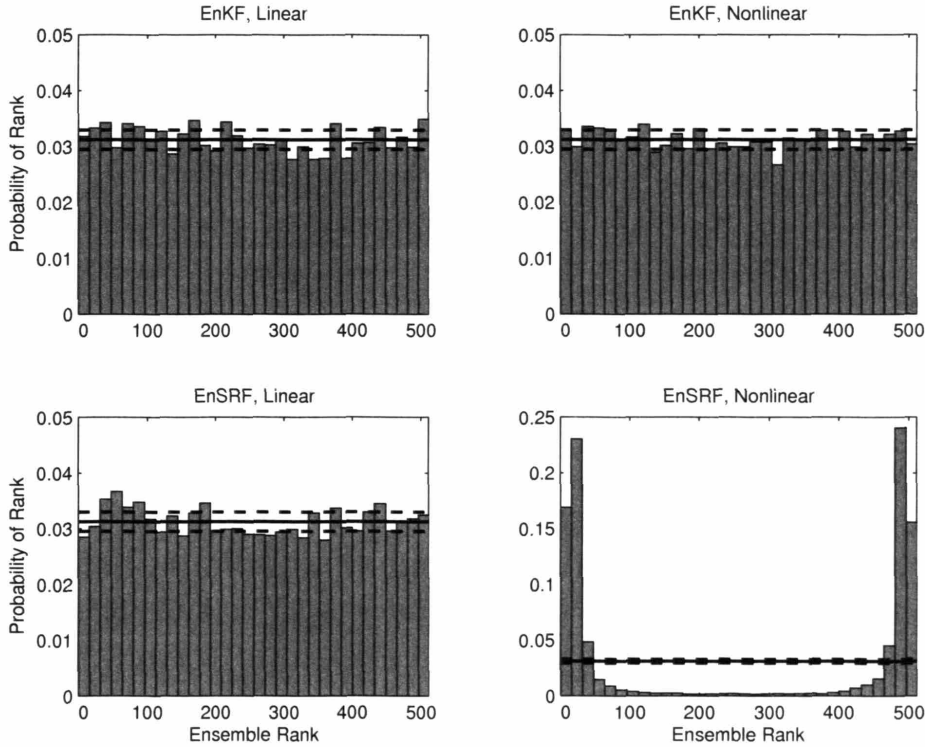


Figure 3-10: Lorenz 95 Ensemble Rank Histograms: shows the same information as figure 3-9 except for the Lorenz 95 system. The linear case has an observation time of 12 hours and an observational uncertainty of  $2.0 \times 10^{-4}$ , and the nonlinear case has an observation time of 12 hours and an observational uncertainty of 0.2.

### 3.4.2 Double-gyre barotropic ocean model

To test consistency in a larger, two-dimensional, spatially extended geophysical model, we examine the filters' behavior over OSSEs with a one layer version of the Harvard Open Ocean Model (Haidvogel et al. 1980; Miller et al. 1983), a quasi-geostrophic model. We run the model with zero buoyancy at the top and bottom boundaries, making the dynamics *de facto* barotropic. Potential vorticity is lost through bottom and lateral friction as well as from numerical filtering. The model is forced by a surface Ekman pumping to give the

classic basin double-gyre western boundary current flow field. The version of the model we used is solved with finite elements and a second order Adams-Bashforth time-stepping routine. The model is well-documented in its original references as well as more recently in an appendix of Moore et al. (2002), interested readers should seek further details there.

The system's state dimension is 2,145 (= one prognostic variable at  $33 \times 65$  gridpoints). For our OSSEs, we use  $N = 511$  ensemble members, but we only observe at 29 locations, all fixed locations chosen to be near dynamically interesting places, either along the western boundary current or at the separation between the gyres. Though it is unphysical, we choose to observe potential vorticity directly so as to minimize complications in the observation operator. Again, we choose to vary observational uncertainty and observation time length to control the degree of linearity. We have swept quickly through the linearity range to find suitable cases to study. We find that errors seem to grow quite linearly using an observational uncertainty that is 5% of a climatological error level and an observation time of 2 days. The slightly nonlinear case we consider here uses three times this observational uncertainty and an observation time of 10 days. The results below consider 10,000 assimilations per OSSE.

A figure of the retrieved rank histograms is omitted because it is largely redundant — it looks indistinguishable from the analogous figure compiled for the Ikeda system, figure 3-9, with roughly uniformly distributed rank histograms for both filters in the linear case and for the stochastic filter in the nonlinear case, and with a U-shaped histogram (peaks in the end bins) for the deterministic filter in the nonlinear case. Whereas above with the simpler models, where we observed every state component at each assimilation time, we saw that the two filters produced essentially indistinguishable ensemble means from one another within the systems' respective linear regimes, we do not find this here with the ocean model. However, it is interesting to note that the filters' average rms analysis error statistics are identical. So, as above, we find that both filters are producing equally accurate updated ensembles for the linear regime. Skewness and kurtosis measures complement these indicators in the expected manner showing that both filters are producing ensembles closely resembling Gaussian PDFs.

In the nonlinear/non-Gaussian case, the rms analysis error statistics do not differ that significantly, the stochastic filter is about 10% more accurate on average over the 10,000 assimilations considered. The skewness and kurtosis time series show that the stochastic filter maintains a very Gaussian ensemble estimate whereas the deterministic filter's ensemble estimates are frequently non-Gaussian, though the measures never approach their maximum values.

*The behavior we observe over the range of linearity for this barotropic ocean model is similar to that which we observed with the Ikeda system and the Lorenz 95 system. We again conclude from our ensemble assessment tools that the filters' behavior in this geophysically relevant model is consistent with the geometric interpretations devised from the simple systems.*

## 3.5 Conclusions

We have sought to understand how and why stochastic and deterministic ensemble-based data assimilation systems can differ in their behavior. From geometric considerations in one and two-dimensional state spaces we have confirmed that both filters behave as expected and perform well in regimes with linear error growth dynamics. In the linear range, the filters are simply combining two Gaussian PDFs, the ensemble distribution and the observations, in the near-optimal fashion theory predicts the EKF will accomplish. In these small state spaces we can visually confirm that the ensemble distributions are in fact Gaussian in shape. Measures of the skewness and kurtosis of the ensembles of course confirm this. By performing OSSEs with the Ikeda system we were able to show that both filters produce updated ensembles that are consistent with the PDF of which truth is a member.

Visualizing ensemble distributions in these small state spaces has also shown how ensemble estimates of truth can be affected by nonlinear error growth dynamics. By allowing errors to grow to moderate amplitudes, initial Gaussian ensembles can become non-Gaussian. When used with sufficiently accurate observations, the random sampling necessary in stochastic filters allows the updated ensemble to assume much of the Gaussian form of the observations, thereby tending to erase the non-Gaussian higher moments nonlinear error growth has generated. These updated ensembles look Gaussian when plotted, and skewness and kurtosis measures confirm they have only small departures from normality. OSSEs run with the Ikeda system show that this tendency for the updated ensembles to repopulate Gaussians is effective at keeping its estimates consistent with truth's PDF, as told by their uniformly distributed ensemble rank histograms.

Deterministic filters, on the other hand, readily maintain non-Gaussian moments in their ensemble distributions by virtue of their direct transformations. We have observed that successive updates of non-Gaussian ensembles can shear the edge ensemble members away from the mean, particularly as they are transformed well off the system attractor. These edge members can easily become far outliers which, under enforcement of having the EKF-predicted analysis error covariance, can lead the main ensemble cluster to tighten. For

cases with fully nonlinear error growth, we observe the updated ensembles to often assume configurations with almost the entire ensemble having collapsed onto one state so as to counter-balance the very few (or even sole) distant outliers. This distribution approaches the limit of maximum ensemble skewness and kurtosis (i.e., maximum non-Gaussianity). The outliers give the data assimilation system the degrees of freedom it needs to match the desired error covariance. Because the ensemble distributions are often skewed and leptokurtic (having positive kurtosis), truth often ranks near the edge members, between the outliers and the main cluster. This is seen to generate U-shaped ensemble rank histograms, indicating that deterministic filter ensemble estimates in nonlinear regimes are not sampling from PDFs that are consistent with truth.

While this geometric interpretation is satisfying, it is based on a contrived one-dimensional example and a two-dimensional chaotic map, and does not necessarily apply to the large and spatially extended models forecasters and geophysicists often care about. Hence, we have extended this interpretation to two higher dimensioned models through the ensemble assessment measures of rms analysis error statistics, ensemble rank histograms, and univariate measures of ensemble skewness and kurtosis. We have found that the retrieved measures in these more complex models are in accordance with the expectations formed from the simple model. These indicate an ability on the stochastic filter's part to handle nonlinearity better than the deterministic filter, or at least over a larger parameter range.

We stress that this is not judgment on a "better" filter, it is simply a consequence of the formulations that comprise the filters. Each individual implementation of an ensemble filter will in a sense determine which aspects of its behavior are important. Also, armed with this insight, something might be able to be done to address these problems. Given that deterministic filters having trouble with nonlinearity seems to spawn from generating non-Gaussian ensemble distributions, one can imagine alleviating these problems in various ways like, for example, periodically resampling the ensemble in the "bootstrapping" fashion suggested by Anderson and Anderson (1999), essentially using a kernel density approach borrowed from particle filtering.

We have purposefully considered only large ensemble sizes so as to avoid some of the "messy" considerations of operational implementation such as ensemble covariance localization and covariance inflation. Even so, we feel the above has begun to shed light on such an understanding, and that the issues raised may become direct concerns for numerical weather prediction in the years to come as computing power allows implementation of very large ensembles. Similarly, we have also restricted ourselves to consideration of the perfect model scenario to hone our focus on the pure implications of the filter formulations. Addressing

model error is a burgeoning area of research and remains a difficult complication to treat. Clearly, model error is too large of a consideration to generalize over except to say that the results depend on the type of error including and how it is compensated for. We have found the inclusion of model error could possibly change the above results and interpretations, depending on how the error is handled. If model error were to be addressed by the inclusion of a model error covariance term, typically written as  $\mathbf{Q}$ , its stochastic nature would tend to erase higher order moments and render subsequent analyses more Gaussian. If  $\mathbf{Q}$  is sufficient to compensate for the model error present and if it can keep the filters' estimates from diverging, then the deterministic filter's ensembles are found to be more reliable than in the above perfect model scenarios. However, if one were to address model error by simply including a large ensemble covariance inflation factor, then the geometric argument is generally found to hold since covariance inflation does not change the structure of ensemble distribution, only the members' distances from the ensemble mean.

The key to handling or avoiding undesirable behavior in these filters is simply understanding them and their limitations. To this, we note that even though we have studied a model hierarchy, we have not approached the dimensionality and complexity of modern day numerical weather prediction models. That these interpretations extend to those systems is likely, but impossible at this point to assert.



# Chapter 4

## Position Errors & Lagrangian Estimation

It seems rather obvious that if a data assimilation method's assumptions are violated, then it will perform differently than expected. In the previous chapter, it was found that the various posings of ensemble-based Kalman filters begin to fail at different points in parameter space, and, moreover, they fail in different ways: neither the stochastic approach nor the deterministic approach is *a priori* better than the other when the prior ensemble is non-Gaussian, they are both wrong. Chapter 2 discussed how these ensemble-based methods seek to approximate the correct approach to probabilistic data assimilation (DA), and impresses that they are the state-of-the-art. At root, it is nonlinear dynamics that allows for nonlinear error growth and the resulting non-Gaussian probability density functions (PDFs, also short for "probability distribution functions," should they be non-differentiable). By neglecting the higher order moments of PDFs, one is discarding potentially important information (e.g., system attractor structure) needed by the DA system to correctly update an estimate. The rest of this thesis concerns whether some of the important discarded information can be reinstilled to the estimation process in certain geophysically relevant situations by consideration of the phenomena at hand.

Chapter 3 considered generic sources of nonlinear error growth, where initial errors are allowed to grow large. The controls over error growth are contained in the observation system, namely the density, frequency, and accuracy of coverage. Given that DA seems to work on the whole in numerical weather prediction (NWP) scenarios, this chapter addresses whether nonlinear error growth is a concern for NWP. *We argue that nonlinear error growth and non-Gaussian PDFs are legitimate concerns in NWP scenarios, and that they spring from a specific source. The main argument is: a) errors arising from mis-positioned coherent*

*features can readily attain non-Gaussian distributions, and b) coherent features are common, if not ubiquitous, in geophysical fluid flows. The main problem arising from coherent features is that their fundamentally Lagrangian errors are treated in an Eulerian sense by the DA systems used for NWP models.*

This chapter has two main sections. The first section argues why coherent features and their associated position errors are relevant to NWP. It gives an overview of previous studies concerning position errors, including the operational approaches taken for tropical cyclone prediction since tropical cyclones serve as a good archetypal atmospheric feature. This first section serves to motivate the rest of the work in the thesis. The second section considers filtering a point vortex model, a model whose state vector is purely Lagrangian in nature. By using a Gaussian vortex approximation for the point vortex model's  $\delta$ -functions, it contrasts ensemble-based DA's ability to update a Lagrangian state vector and its corresponding Eulerian state vector, obtained from plotting its fluid fields on a grid. The experiments presented show that there are cases where an ensemble Kalman filter (EnKF) cannot update the gridpoint versions of the point vortex states even though it is able to update the Lagrangian position representation. The main difference between this chapter and the two following it is that the DA methods in this chapter generally work within a Lagrangian framework, and hence are of mainly academic interest, whereas the following chapters consider DA methods more appropriate to the gridpoint based models that comprise NWP.

## 4.1 Features and position errors

Geophysical fluid dynamics, of which NWP is a subset, is nonlinear, and therefore capable of generating non-Gaussian error PDFs. Given the results from the previous chapter, namely that state-of-the-art DA systems fail when faced with non-Gaussian PDFs in relatively simple models, the question naturally arises: should one worry about the state-of-the-art methods (and currently implemented methods) failing when applied within NWP? As was mentioned in the previous chapter, by and large, ensemble-based filtering has proven effective in many small and medium-scale sized problems. However, as the previous chapter made clear, the success of maintaining a regime of linear error growth depends critically on the accuracy, density, and frequency of available observations. For large-scale baroclinic waves, the current observational network seems adequate for constraining error evolution; however, as evidenced by some of the difficulties that arise within tropical cyclone forecasting (e.g., Zou and Xiao 2000; Aberson 2002) and mesoscale forecasting (e.g., Zhang et al. 2002;

Zhang 2004), there are clearly some scales of motion that are still not well constrained by the observations. Many of these scales are associated with strong coherent features, often bearers of severe weather. Some of these features are large-scale enough to be resolved by NWP models, and it is these features that constitute a likely source of non-Gaussian error PDFs. As we shall see, the non-Gaussianity generally arises from having position errors present, also called alignment errors.

#### 4.1.1 Position errors

Many classes of problems and phenomena that scientists try to predict exhibit errors of alignment, where some coherent feature is improperly positioned yet properly, or at least reasonably, shaped. Such errors are ubiquitous in Lagrangian control systems, where problems like guiding robotic arms and landing aircraft are common. NWP is no exception. *Even though most state estimates of the atmosphere are fundamentally Eulerian, being the output from large, discretized numerical models, the phenomena whose evolution they approximate have apt Lagrangian fluid descriptions.* While technically, Lagrangian fluid dynamics describes fluid *parcel* evolution, it is a useful construct as well for describing spatially-extended coherent features. While many scientists have tried to give a formal definition to “feature” in the context of fluid dynamics, we resist doing so here. Features are often noted for being *coherent*. For example, Provenzale (1999) defines a coherent vortex as a region of circulating flow that keeps its identity for times much longer than the local eddy turnover time, which in turn can be defined in terms of the local average enstrophy. Features are often noted for being tight or compact, and this is often synonymous with being *strong* in some sense: this identification is typically associated with tight gradients or “bunching of contours” when displayed on a map. Both notions, coherent and strong, approach the featuredness addressed in this thesis. Without being too much more specific, we generally treat a feature as something that can be identified (i.e., pointed to) on a map. In fact, it is the very notion of geographic location, or position, that matters most.

Strong, coherent features appear quite readily in fluid flows, especially rotating flows. Phenomenologically, this includes fronts, jets, convective storms, solitary waves, and coherent vortices like tropical cyclones, Gulf Stream rings, so-called “Meddies,” and even larger-scale features like cut-off lows, blocking highs, and the polar vortex. Perhaps one of the most prevalently studied examples of feature emergence is that from so-called “geostrophic turbulence” — rotating, barotropic, high Reynolds number flows typically initialized by fields of a given energy spectrum with randomized phases. Many scientists have considered this problem, though it was McWilliams (1984) who first concentrated on the coherent vortices

that commonly emerge. Many other studies have since re-confirmed this behavior (e.g., McWilliams 1990; Provenzale 1999) and found coherent vortices spontaneously forming at length scales different from those dominant in the initial energy spectra. The formation is often attributed to the inverse cascade of energy (i.e., from smaller scales to larger scales) in two-dimensional turbulence, but as pointed out by Provenzale and Balmforth (1998), this is not sufficient for explaining the vortices' emergence. Another line of feature formation studies has followed the instabilities forming from unstable jets. A very thorough study of the parameter space of a barotropically unstable jet was carried out by Flierl et al. (1987), and they found the jet consistently broke into coherent vorticity patches whenever  $\beta$  was small enough. The long line of studies on the emergence of baroclinic instabilities and frontal cyclones could be considered studies of feature generation as well, depending on one's notion of a feature. Given that many features seem to emerge from wave-like phenomena, there is not always a clear dividing line between a general fluid field and an embedded coherent feature — it is for this reason that we resist strict definition, however, we suggest that the existence of closed contours helps delineate features.

While day-to-day aspects of NWP are mainly concerned with the propagation and evolution of baroclinic waves and their surface highs and lows, it is the often severe weather associated with strong features that the public most strongly depends on forecasters to handle correctly. This is for the obvious reasons concerning safety and property. Given that these crucial forecasting situations typically concern strong, coherent features marked by position information, one inevitably questions whether NWP can handle and correct position errors. In general, one finds that implementable DA methods are not able to handle and correct position errors because position errors violate their basic assumptions. As we show in the next chapter, when a feature's position error is comparable to its own characteristic length scale, then the gridpoint-based error distribution can be shown to be non-Gaussian, even when the position errors themselves are Gaussian. Moreover, we also demonstrate in the next chapter that the nonlinear error growth leading to the non-normal error distribution can easily violate the tangent linear approximation. Traditional (yet still "state-of-the-art") DA systems will produce poor analyses in these situations leading in turn to poor forecasts, precisely in the situations where good forecasts are depended upon.

The details of the above claims and assertions will be given in the next chapter, but for now we offer a heuristic example to highlight the main idea, shown in figure 4-1. The dashed black line in the top panel shows a hypothetical one-dimensional base profile with a "bump" in the center of the domain represented over 200 gridpoints (denoted  $s_i$ ,  $i = 1, 200$ ), the profile's actual form will be discussed in the next chapter. The thin blue lines denote

a 50 member ensemble that has been generated by simply shifting the base profile left or right by a Gaussian random amount. The dark blue dotted line shows the mean of the ensemble members (one dot per gridpoint). Note that the mean looks different than any individual member in the region of bump discrepancy. If we call the base profile the expected truth state, then the deviations between it and each ensemble member are the *errors* of each ensemble member. If we treat the errors as additive errors like estimation methods based on the Kalman filter often do, then the middle panel shows the ensemble of these errors, found by subtracting each member from truth. The dark blue dotted line shows the ensemble mean error. Note the striking coherency in the light blue lines, indicating their non-Gaussianity. Note also that their mean is not a zero vector, indicating that the errors are biased. To further demonstrate their non-Gaussianity, the bottom panel shows a normalized histogram of the value of each ensemble member’s additive error at a gridpoint within the region of bump discrepancy ( $s = -3$ , the 85th gridpoint, marked by the vertical dotted line in the middle panel), compiled from a similarly constructed 1000 member ensemble. Note the obvious non-Gaussian structure. When treated purely additively, these errors are non-Gaussian and biased, even though we know each is the result of a single unbiased Gaussian random number.

In the rest of this section, we first review forecasters’ previous acknowledgments of position errors and their associated forecasting difficulties. There is a long history to position errors because they are one of the natural components of a qualitative forecast. We then review previous studies that have attempted to handle, classify, and correct position errors. We begin this section with a discussion of two approaches that have specifically looked at addressing features in state estimation, “contour analysis” and “feature calibration and alignment,” though the latter is more germane to the next two chapters. The section concludes with a discussion of tropical cyclones since the tropical cyclone could be considered the archetypal atmospheric feature. Also, the historical development of tropical cyclone prediction efforts punctuates the troubles that features can bring to true operational NWP.

### 4.1.2 Historical treatment of position errors

Forecasters have long been aware of position and storm arrival timing errors. Even before the days of widespread transportation, it was known that storms traveled, and Benjamin Franklin is generally credited with being the first to recognize, in 1743, that storms need not propagate in the direction their surface winds blow (e.g., Lyons 1997). When distant communication had increased widely enough to begin constructing synoptic weather maps, weather forecasting followed closely behind, though initial attempts were mostly done

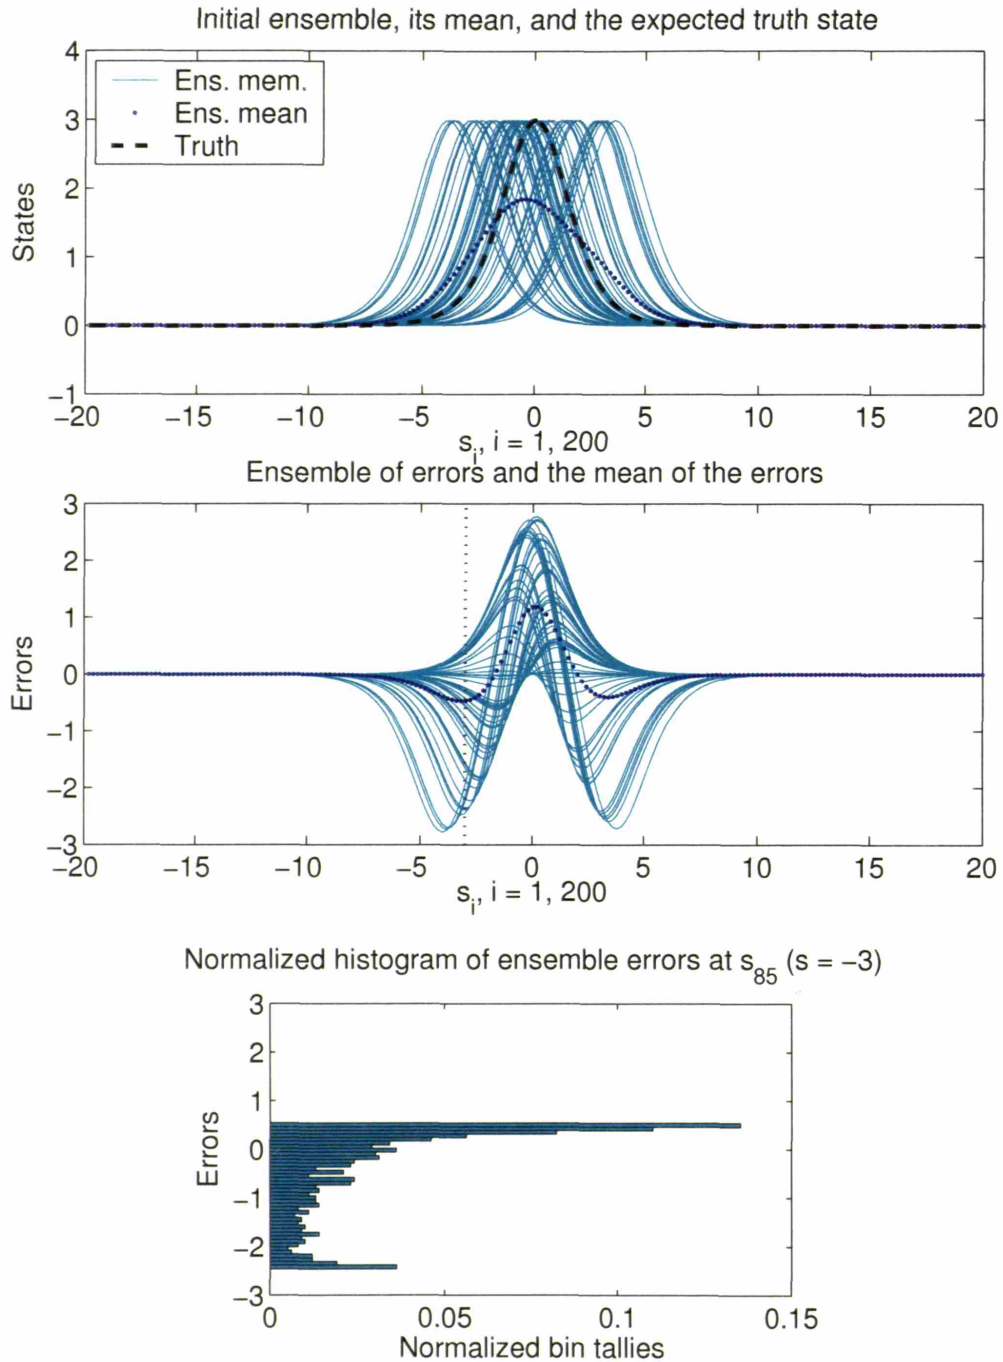


Figure 4-1: The top panel shows a 50 member ensemble of the same “bump” profile shifted a Gaussian random amount to the left or right. The middle panel shows the resulting ensemble of additive errors when the ensemble in the top panel is subtracted from the black dashed truth profile. The bottom panel shows a normalized histogram of the total additive errors at  $s_{85}$ , (i.e.,  $s = -3.0$ ) compiled from a similarly constructed 1000 member ensemble.

by intuition. From the beginning, the importance of position errors was recognized, simply because a qualitative forecast communicates mostly position information anyway. For instance, Henry (1916) stresses the importance of correctly positioning the “dominating pressure control” because of the large forecast errors that can arise from small positioning errors.

As NWP developed in the computer-age, and as objective methods of DA were devised, scientists again began to address position errors. Early methods of DA were quite simple (e.g., Panofsky 1949; Gilchrist and Cressman 1954; Eliassen 1954; Barnes 1964), and did not specifically address position errors. It was not until the exploratory study of Charney et al. (1969) that meteorologists sought what has come to be termed “four-dimensional” data assimilation. Charney et al. had in mind using satellite observations to constrain the whole atmospheric state, even over successive times. As others rose to their challenge, it became clear that a significant source of error was so-called “phase error” (i.e., timing error). Blumen (1975) studied the causes of phase error and its contribution to total error under the premise that simple truncation error from coarse model discretizations caused the propagation speeds of atmospheric features to be under-estimated by models. The position and phase errors considered most often in these early studies were associated with mid-latitude cyclones. As model resolution has increased and the observational network improved, position errors of mid-latitude cyclones have generally decreased to a point where they are typically much smaller than their length scale (the Rossby deformation scale), at least over well-observed land masses. McMurdie and Mass (2004) have diagnosed that rather spectacular short-term forecast failures of storms coming off the Pacific Ocean are attributable, at least in part, to errors in the forecasts’ initial storm positions, though those position errors are still quite small compared to the characteristic cyclone length scale. Mid-latitude frontal cyclones aside, many of the aforementioned strong, coherent features (e.g., tropical cyclones) still exist in NWP and can give problems to prediction efforts.

Though not motivated by NWP, another strain of feature oriented meteorological research has been in the field of model validation, particularly in regards to a model’s representation of storm tracks. The first identifications of storm tracks were feature-based (as cited in Hoskins and Hodges 2002), though storm tracks have come to be defined by Eulerian, gridpoint-based statistical analyses of eddy energy (e.g., Blackmon et al. 1977). Various posings of the feature-based approach have been made. Hodges has been a recent proponent of applying pattern recognition techniques to general circulation model (GCM) output and satellite images (Hodges 1994, 1995, 1998, 1999; Hoskins and Hodges 2002; Anderson et al. 2003; Delsol and Hodges 2003; Delsol 2005). He and his coauthors have of-

ferred automated procedures for both feature identification procedures and feature tracking procedures. Much of their work can be considered characterizing climatological position error.

### 4.1.3 Contour analysis

One of the first studies of a systematic approach to handling position errors is by Mariano (1990). Mariano introduced the idea of contour analysis for “melding geophysical fields.” Mariano had in mind the problems that occur when one has two completely specified estimates of the same state that have a sizable discrepancy in position between one or more coherent features. He showed through schematic diagrams and examples that simply averaging the two fields yields an analysis with smeared, weakened features that do not resemble the features in either of the contributing estimates. Mariano sensed that the typical process of forming an analysis was not well suited to estimates with position discrepancies, and correctly diagnosed that it was because traditional methods treat errors in an Eulerian sense while the errors at hand are Lagrangian in nature. He offered a solution to this problem where the initial estimates are first contoured in one or more fields, that is, diagnosed for the isopleths of one or more quantities. Once the contours are known, one then draws a correspondence of isopleths between the two estimates. Corresponding isopleths are then divided along their arclengths into an equal number of segments, and the spatial coordinates of the segment nodes are found through cubic spline interpolation. The analysis is formed by finding a weighted average of the segment node positions and interpolating the new field to find the gridpoint values of interest. Figure 4-2 is adapted from a figure in Mariano’s original paper schematically showing how contour analysis works. The top two panels each show the streamfunction of a hypothetical geophysical field. Each field contains a front and a vortex, but the features are differently positioned. The lower left panel shows the simple gridpoint-by-gridpoint mean, that is, the “traditional analysis,” and the lower right panel shows Mariano’s contour analysis. Note the smearing present in the traditional analysis compared to the contour analysis.

Contour analysis is a geometrically attractive approach to data assimilation within physical space (i.e., discretized model space). It directly addresses position errors in a Lagrangian sense, and the resulting analysis is not smeared or weakened as it was before. Extended to an ensemble context, one can imagine how the “contour analysis ensemble mean” of the initial ensemble in the above schematic example (figure 4-1) would more closely resemble truth than the gridpoint-based ensemble mean. However, there are several problems to overcome before contour analysis could become a regularly used DA method. First, one must decide

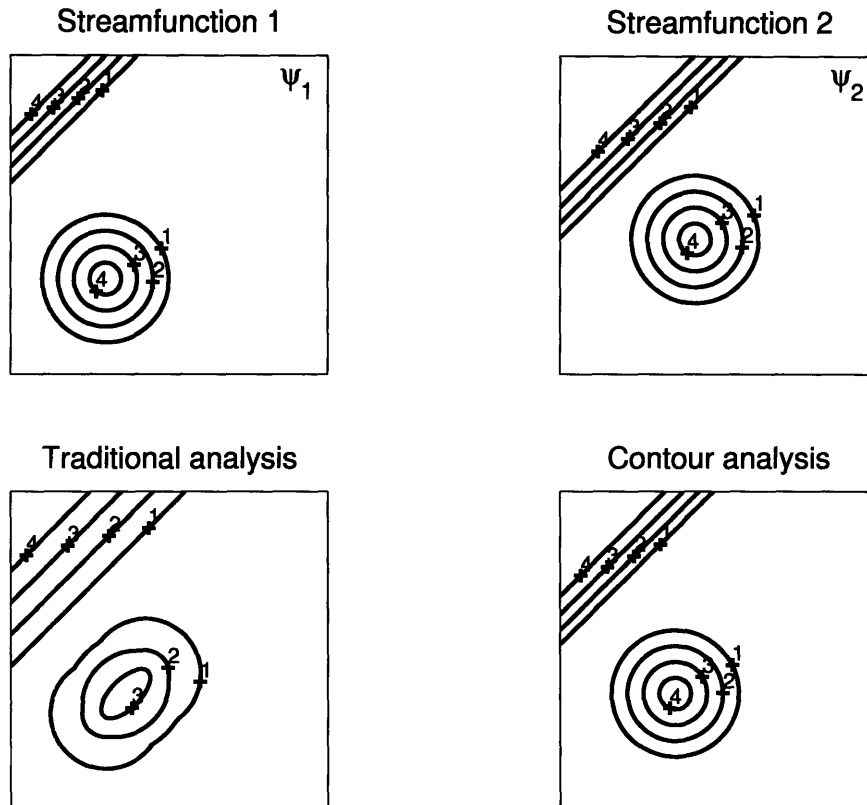


Figure 4-2: The top two panels show two different nondimensional streamfunction fields. The bottom left panel shows the traditional analysis obtained by averaging each gridpoint value. The bottom right panel shows the contour analysis obtained by average contour coordinates. (Adapted from Mariano 1990)

how to handle the analysis when one state is not completely specified. Observations rarely cover the entire state, so in order to access their information content one needs to overcome a difficult, if not ill-posed, correspondence problem. Second, one must decide on a strategy for situations where a one-to-one correspondence cannot be established between features in the available estimates (including the observations). In his paper, Mariano addresses the one-to-one correspondence problem and offers several suggestions to help renormalize the fields. Quite judiciously, he concludes that if one cannot massage the fields into rough correspondence, then one should not attempt contour analysis since the fields differ too widely. One last concern with contour analysis is that it is unclear how estimates' uncertainties posed in an Eulerian sense will translate to a Lagrangian sense for use in the weighting of contour node positions. Interestingly though, Mariano and Chin (1996) go further to offer an approach to study the variability of a system within the coordinates defined by contour

analysis, what they call “empirical orthogonal contours.” Regardless of any problems, contour analysis stands as the first systematic approach to explicitly treating position errors for the purposes of state estimation.

### **Contour analysis within contour dynamics models**

Though Mariano seemed to have the application of melding geophysical fields in mind when he devised contour analysis, it seems it is also well suited for a particular brand of numerical evolution called contour dynamics models. Most contour dynamics models assume that a two-dimensional Euler fluid (i.e., inviscid and incompressible — Euler fluids will be discussed further in the next section) is divided into regions of uniform fluid vorticity whose boundaries are denoted by contours. The continuous contours are represented by a finite number of nodes, quite like Mariano sought to discretize curves in his contour analysis scheme. Since an Euler fluid’s evolution is determined by its vorticity distribution, if the assumption of piece-wise uniformly distributed vorticity is reasonable, then one can represent the nodes on the contours as point vortices (i.e.,  $\delta$ -functions in vorticity) of a given strength. Hence, a contour dynamics model is simply a carefully initiated point vortex model whose vortices of like strength are understood to be connected. The point vortex model will be considered in more detail in the next section.

A problem frequently encountered in contour dynamics models is that as the nodes evolve, the resolution along the curves quickly becomes non-uniform. Though the strength of each initial vortex is conserved throughout integration, there is no such constraint on a contour’s arclength. As the length expands and contracts through processes like filamentation, the nodes on the contour either spread widely apart or become essentially collocated. This Lagrangian advection can get quite complicated to track, owing to why contour dynamics models are often used in relatively quiescent regions like the stratosphere (e.g., Polvani and Plumb 1992) or only over relatively short time periods (e.g., Nolan 2004). This resolution issue has been partly solved by a process called “contour surgery” (Dritschel 1988), where a diagnostic process periodically re-distributes nodes to ensure more uniform resolution. While this has proven effective in evolving contours, the fact that the number of nodes and their positions are periodically changed (i.e., a variable-length state vector) implies that their state estimation by traditional means would be a complicated feat indeed. Contour analysis, however, is naturally posed to resegment contours regardless of the number of nodes and their locations after integration.

#### 4.1.4 Feature calibration and alignment

There is a growing literature under the broad categorization of what has become known as “feature calibration and alignment” (FCA). Calibration roughly refers to correcting a feature’s amplitude while alignment is reserved for the feature’s position (or phase, as it is sometimes called). These studies have grown from recognizing that Eulerian analysis and forecasting approaches are poorly suited to handling Lagrangian errors. Hence, these studies are closely aligned with ours here. As such, they will be considered in more depth in the next chapter. For now, we highlight the important results.

The FCA literature began when Hoffman et al. (1995) introduced a formalism for diagnosing so-called “errors of adjustment,” which is a decomposition of errors into alignment errors, calibration errors, and a residual. Hoffman and Grassotti (1996) applied these ideas to data assimilation with satellite data, Grassotti et al. (1999) revisited that problem with the addition of radar data, and Nehrkorn et al. (2003) delved deeper into making the necessary decompositions and procedures more objective. Dickinson and Brown (1996) applied a nonlinear interpolation procedure to displace storm representations to synoptic time, and Alexander et al. (1998) applied standard image-morphing techniques to model estimates to make them better agree with position and shape information taken from satellite images. Brewster (2003a; 2003b) examined techniques of correcting alignment errors within mesoscale applications. His approach consists of first exhaustively searching for alignment errors (what he calls “phase errors”) as the minimum of a cost function evaluated from the weighted mean-square difference of a smoothed forecast field and dense radar observations, and then applying shift corrections to his forecast field either at the analysis time or over an assimilation window.

In contrast with this literature, the work within this thesis is aimed at implementable estimation strategies for use in *ensemble forecasting scenarios* when relatively sparse observations are available, though we do return to the error decompositions that Hoffman et al. pioneered. We aim to capitalize on applying FCA-like ideas to compensate for the potentially harmful assumptions of Gaussian error PDFs.

We next describe operational prediction efforts for handling position errors in the realm of tropical cyclone prediction.

#### 4.1.5 Tropical cyclone prediction

A good candidate for the archetypal atmospheric feature is the tropical cyclone, and as such its treatment demonstrates the typical problems associated with predicting coherent

features. Indeed, to circumvent these problems, the tropical cyclone community has found need to design and work with bogus vortices. This subsection briefly summarizes operational efforts to track and predict tropical cyclones, it then describes the use of bogus vortices in tropical cyclone prediction, and it concludes with a discussion of ensemble prediction efforts to date aimed at tropical cyclone prediction. This is not a thesis on tropical cyclone prediction *per se*, but tropical cyclone prediction is a likely field to benefit from our work.

### **Track forecasting**

Tropical cyclones are strong disturbances, generally compact in length scale, and relatively long lived. Tropical cyclones are typically centered about an eye, and the surrounding circulations are among the strongest in the storm. In a radial sense, the average tropical cyclone profile has its maximum wind strength close to the eyewall, and the winds, on average, decrease linearly toward the eye, and as  $r^{-1/2}$  outward from the storm. Tropical cyclones generally depart from perfect axisymmetry, though not too severely in the wind and pressure fields; some skillful models even assume that storms are axisymmetric (e.g., Emanuel 1999). Because of their destructive potential, tropical cyclones are typically well-monitored features. A large part of tropical cyclone monitoring is determining its track, and all gathered data is later used to determine the published “best track,” a least squares smoothing estimate of where the storm is most likely to have been at each time. The NHC’s current statistics on tropical cyclone track forecast errors have 12 hour track errors of order 90 km (DeMaria 1997) and 24 hour track errors of order 160 km (Landsea and Lawrence 2004). These errors are comparable to or greater than the typical scale of a tropical cyclone, or at least its core where wind and pressure variations are greatest. Though predictions have been improving lately (e.g., Kurihara et al. 1998), it is believed that further improvements are possible and necessary (Leslie et al. 1998a).

### **Bogusing**

Contributing to making tropical cyclone prediction a difficult problem are the paucity of observations over the tropical oceans and a lack of a full understanding of the governing physics and dynamics of the tropics (Elsberry 1995). While the latter certainly contributes to model error, the former poses a serious problem to implementable DA schemes where, unconstrained by dense observations, prediction errors can grow very large over short times. To cope with these shortcomings, forecasters have had to resort to inserting a bogus vortices into their initial model states. This practice of *bogusing* is a somewhat unfortunate one to have to perform, but NWP centers that make tropical cyclone forecasts find they need it. In

order to explicitly resolve the structure of a tropical cyclone, model resolution of 5 to 10 km would be needed (Elsberry 1995). At best, typical NWP models today can only resolve the outer circulation of the storm. In order to overcome the lack of observations and relatively coarse resolution of the prediction models, almost all operational tropical cyclone prediction systems are forced to “trick” their DA systems by generating and including synthetic or bogus observations of a given storm’s structure (Leslie and Holland 1995). Often this bogusing is implemented by first identifying and removing the (poorly) observed tropical cyclone signal from the background field and then reinserting a well-observed, “synthetic” vortex at the observed location in the analysis (e.g., Surgi et al. 1998; Kurihara et al. 1998; Heming and Radford 1998). The initial use of bogusing used symmetric synthetic vortices, but recent implementations have switched to asymmetric vortices and “spun-up” vortices so that the vortex flow agrees better with the background state (e.g., Liu et al. 1997). Such studies eventually led to the notion of combining bogusing with DA. For example, Zou and Xiao (2000) used sea-level pressure “observations” from a bogus vortex within an adjoint method DA scheme (see equation (2.46)) applied in a high resolution mesoscale model (MM5) simulation of Hurricane Felix. They found a great improvement over the track forecasts that had been initially issued for the storm, though they note they had not fully eliminated spin-up issues.

While this practice of bogusing has led to improvements in operational tropical cyclone track prediction, it is ultimately an undesirable process to have to perform. Aside from the arbitrariness of defining the vortex signal to be removed and reinserted, bogusing addresses the symptoms of a bad forecast rather than its cause, that is, whatever error led to the bad track forecast in the first place (presumably in the environmental background flow or in the model itself) is left uncorrected. It should be noted that recent experiments in tropical cyclone DA and prediction have managed to avoid bogusing when high resolution satellite derived winds were used (Leslie et al. 1998b; LeMarshall 1998). This is to be expected as very spatially dense observations should alleviate many of the problems experienced in NWP. Hence, the availability of new data and the ability to use it may soon eliminate bogusing, however, correcting large-scale background flow errors will still need addressing.

### **Ensemble prediction efforts**

There have been some attempts at ensemble forecasting for tropical cyclones. As discussed by Weber (2003), there have been mainly two approaches, statistical and dynamical. From the statistical approaches, many have followed the style of Leslie and Fraedrich (1990) who found an optimal combination of a full tropical cyclone forecast model and CLIPER,

the standard climatology-persistence forecast scheme used in tropical cyclone prediction. They found their combined forecast out-performed both the forecasts which formed it and two other commonly used models at the time. Improving upon this, Goerss (2000) found that averaging track position forecasts from three different models yielded more accurate forecasts than any single particular model. Krishnamurti et al. (2000) found similar results with their so-called “superensemble” approach where forecasts from seven different models are regressed toward observations to yield an optimal model combination of multimodel forecasts. The superensemble out-performs the ensemble mean, showing statistical post-processing can help improve over Goerss’s method. Weber (2003) takes a similar approach to the superensemble, except he chooses to weight each model based on its verified performance from the prior year based on certain storm characteristics like structure, location, and motion.

The dynamical ensemble approaches are more in the spirit of the Monte Carlo methods discussed in previous chapters. Some perturbations are produced and added to a current storm field to yield an ensemble of forecasts. None have yet cycled the ensemble dispersion information back into the data assimilation step, but studies such as Puri et al. (2001), where the authors perturbed an initial field with so-called “moist singular vectors,” show that Monte Carlo methods can be effective at spanning the possible position errors that tropical cyclones might incur.

#### **4.1.6 Position error conclusion**

We have argued that errors in the positions of coherent features can lead to non-Gaussian uncertainty. To this point we have only offered the reader the heuristic example shown in figure 4-1 and evidence from other authors’ studies, but we will consider this process in detail in the next chapter when we examine how Eulerian DA systems treat position errors. We have also reviewed previous efforts to address position errors. Phenomonologically, position errors spawning from the prediction of coherent features can be real concerns, and they have been problematic for forecasters and data analysts trouble alike.

In the next two chapters we examine ensemble-based DA systems that address position error within an Eulerian framework. However, before addressing this more difficult problem, we find it instructive to first examine ensemble-based DA systems in a purely Lagrangian framework; this is the concern of the rest of this chapter. As will be made clear in the following chapters, the difference between these two scenarios is a matter of how Lagrangian information is related to the state vector — position is a nonlinear function of an Eulerian state vector, whereas position defines a Lagrangian state vector. Hence, with adequate ob-

servations, we expect that a DA scheme based on linear(ized) operators should perform well in a completely Lagrangian scenario. Much of NWP’s infrastructure is based on discretized, Eulerian models, and it is unrealistic to expect any directly Lagrangian DA systems could ever be useful to the community. However, before considering how Lagrangian errors affect Eulerian-based models and DA schemes, it is worth considering, by way of motivation, how Lagrangian errors affect Lagrangian-based models and DA schemes: at the very least, it will help clarify characteristics we would like the Eulerian DA systems considered in the subsequent chapters to exhibit.

## 4.2 The point vortex model & Lagrangian estimation

Because position errors are naturally posed as a Lagrangian idea, correcting position errors naturally suggests a Lagrangian data assimilation system. This section considers methods of how one might choose to implement an ensemble-based DA system that directly uses Lagrangian information or even works fully within a Lagrangian framework. We first consider a point vortex model, whose state vector is naturally comprised of Lagrangian position information. As expected, ensemble-based Kalman filtering is shown to work quite well within reasonable bounds (i.e., with sufficient observational coverage and using non-pathological first guess fields). This is true even when used with (nonlinearly related) observations of Eulerian quantities like fluid velocity, with two caveats first noted by Ide and Ghil (1997a) (discussed below). The point vortex model proves a useful model with which to study the effects of position errors. In particular, we attempt to simulate the point vortex model’s behavior within a discretized barotropic model. Since point vortices are singularities in vorticity, they need to be approximated in the barotropic model by distributed vorticity patches. One can then vary the areal coverage of the patches relative to their ensemble spread to gauge when position errors become problematic for implementable DA systems — clearly, as the patches increase in size, one expects DA to be more successful.

### 4.2.1 The point vortex model

Although it may seem not strictly relevant to current NWP pursuits, we first consider a dynamical system whose state vector is comprised of Lagrangian variables, a point vortex model. The point vortex model, as posed here, is a specific (singular) solution to a simplified form of the governing equations, the Euler equations (incompressible and inviscid), and therefore is fluid dynamical in nature. Point vortices themselves are  $\delta$ -functions in vorticity, meaning their velocity profiles decrease radially from a singularity at the center as  $r^{-1}$ . The

center singularity notwithstanding, a point vortex is a reasonable first order approximation of atmospheric vortices like tropical cyclones, hence the model supports features. In fact, the model is *defined* by features.

Beginning with the two-dimensional Euler equations, one can derive an expression for the conservation of fluid vorticity, the barotropic vorticity equation:

$$\frac{d\zeta}{dt} = 0, \quad (4.1)$$

where  $\zeta$  is the vertical component of the vorticity vector. Since an Euler fluid is incompressible, its velocity field is non-divergent. This allows construction of a streamfunction, and in two-dimensions, this determines a Laplacian relationship between the vorticity and the streamfunction:  $\zeta = \nabla^2\psi$ , where  $u = -\frac{\partial\psi}{\partial y}$  and  $v = \frac{\partial\psi}{\partial x}$ .

A point vortex model defines the vorticity field as a summation of  $N_v$   $\delta$ -functions in the two-dimensional plane, each with its own position,  $\mathbf{x}_i = (x_i, y_i)$ , and circulation  $\Gamma_i$ :

$$\zeta(\mathbf{x}) = \sum_{i=1}^{N_v} \Gamma_i \delta(\mathbf{x} - \mathbf{x}_i). \quad (4.2)$$

One can invert the Laplacian relationship via the Green's function to find the streamfunction field:

$$\psi(\mathbf{x}) = \sum_{i=1}^{N_v} \frac{\Gamma_i}{2\pi} \log |\mathbf{x} - \mathbf{x}_i|, \quad (4.3)$$

and in turn this can be differentiated to find the components of the velocity field:

$$u(\mathbf{x}) = - \sum_{i=1}^{N_v} \frac{\Gamma_i}{2\pi} \frac{y - y_i}{\|\mathbf{x} - \mathbf{x}_i\|^2} \quad (4.4)$$

$$v(\mathbf{x}) = \sum_{i=1}^{N_v} \frac{\Gamma_i}{2\pi} \frac{x - x_i}{\|\mathbf{x} - \mathbf{x}_i\|^2}, \quad (4.5)$$

where  $\|\cdot\|$  denotes the  $\ell_2$  norm (i.e., magnitude) of its vector argument. Note the singularities of the fields at the positions of the point vortices. Because the vorticity field is comprised only of  $\delta$ -functions, their positions and strengths characterize the entire flow field. That is to say, the continuous PDE solution reduces to following the evolution of three variables per point vortex, and in fact, since the circulations for each point vortex remain constant unless further dynamics are included, the solutions are effectively characterized by two variables per point vortex, the  $x$  and  $y$  position. If, say,  $N_v = 4$ , then the dimension of

the state vector is 8. The relevant dynamics are simple mutual advection, the other three vortices determine the evolution of the fourth. Once the positions (and circulations) are known, then the entire fields of  $\zeta$ ,  $\psi$ ,  $u$ , and  $v$  are specified. The point vortex model is thus a reduction of the full PDEs to  $2N_v$  ODEs for the positions of each point vortex in the domain, where  $u(x_i, y_i) = \frac{dx_i}{dt}$  and  $v(x_i, y_i) = \frac{dy_i}{dt}$ .

An Euler fluid is, by definition, frictionless, and therefore, the system has no attractor. In fact, as reviewed by Aref (1983), a point vortex model is a Hamiltonian system. For collections of point vortices in an unbounded domain, the system dynamics remain *integrable* up to  $N_v = 3$ , meaning that the solution can be expressed as an explicit function of time in terms of integrals of expressible functions. Integrable solutions are, at worst, quasi-periodic (Aref 1983). For  $N_v \geq 4$ , the system dynamics can become nonintegrable in general and must be approached numerically; however, certain configurations can still be integrable, as when four point vortices of equal strength denote the vertices of a square in the two-dimensional plane. Nonintegrable Hamiltonian dynamics is qualitatively similar to chaotic dynamics in that both show a marked sensitivity on their initial conditions, but only chaotic dynamics have an attractor (due to the dissipation). Some have referred to state space portraits of the obtainable states of nonintegrable Hamiltonian dynamics as “chaotic seas”.<sup>1</sup> Nonintegrable Hamiltonian dynamics can be aperiodic, and look chaotic. An example of this, a variant of an example considered in Aref and Pomphrey (1982), is included in figure 4-3. The left panel shows the initial positions of four equally strong point vortices. The initial positions are perturbed slightly from the vertices of a regular rectangle with an aspect ratio of 2. The right panel shows the ensuing trajectory *in physical space* (as opposed to state space) of the upper left-hand point vortex. Note the irregular path traveled, and also note the seemingly bounded size of the orbits. Nothing *per se* is constraining the point vortices to the area except for the internal advective dynamics — if one of the vortices had a strength equal but opposite in sign, then it is very likely it would find itself in a dipole pair and propagate off toward infinity. Changing one of the initial positions slightly leads to an (on average) exponentially diverging trajectory from the original one that nonlinearly saturates at the “climatological error level.” Also, the effective degrees of freedom of the system are fewer than the formal degrees of freedom,  $2N_v$ , because the dynamics effectively work only with the vortex separations, not their absolute positions.

Point vortex models have a Lagrangian state vector comprised of the coordinates where  $\delta$ -functions in vorticity are located. This is a point-wise notion. At first glance, this may

---

<sup>1</sup>Chaotic seas are typically reflected in Poincaré sections since the dimension of the system is often too high to fully visualize

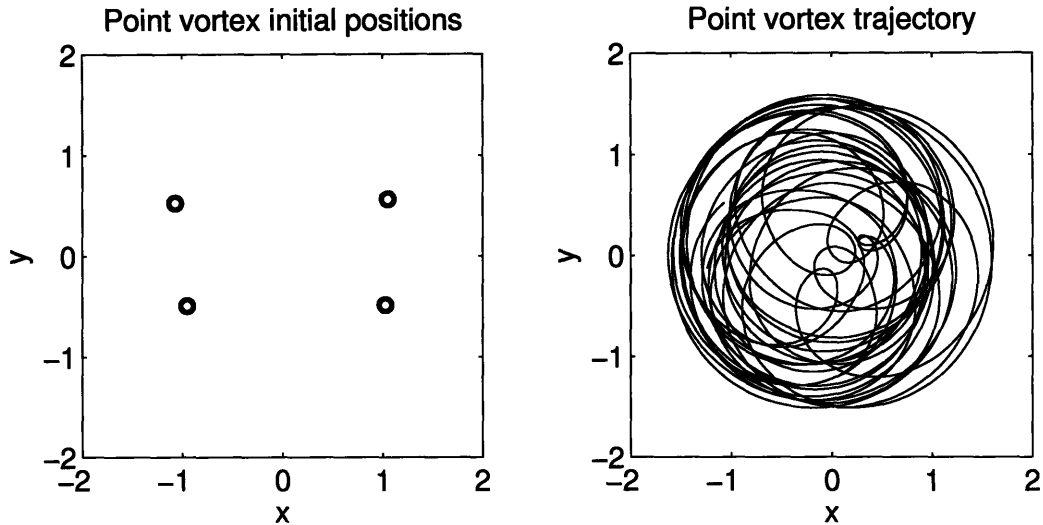


Figure 4-3: The left panel shows the initial point vortex conditions, and the right panel shows the ensuing trajectory of the upper left vortex.

seem restrictive, for real fluid dynamics is on the continuum, not point-wise. There is a literature on so-called vortex methods where scientists try to simulate the full equations of fluid motion (even with including viscosity) by evolving the positions of a large number of point vortices (e.g., Hald 1979; Aref 1983). These methods can be shown to converge to the correct equations as the number of vortices  $\rightarrow \infty$ , so the point vortex model is perhaps not quite as restrictive as it may seem. However, as posed here, the point vortex model is evolving the positions of point vortices, not approximating a fluid continuum. As mentioned above in section 4.1.3, one common intermediate model between evolving point-wise features and the full continuum is a contour dynamics model.

Given that a point vortex model reduces the full Euler equations to a system of  $2N_v$  ODEs, it can be considered, in the spirit of Bjerknæs, as “just another dynamical system.” From a mathematical point of view, a state can be evolved, observed, and estimated, irrespective of its geophysical meaning. However, since the point vortex model has a fluid dynamical interpretation and supports coherent features, it is an instructive beginning to see how DA methods handle it.

## 4.2.2 Previous point vortex predictability studies

Point vortex models have been studied thoroughly for well over a century (e.g., Aref 1983). While discovery and discussion of their nonintegrability for given vortex configurations

certainly has predictability implications, it seems that studies of their predictability were not considered until relatively recently. For instance, it is difficult to find studies where scientists performed state estimation in an effort to improve forecasts. Coming closest, in a rather academic study of the four equal strength point vortex model, Novikov and Sedov (1978) conclude with a suggestion that weather prediction may be so difficult because of nonintegrable vortex interactions through cyclones.

It was Ide and Ghil (1997a) who first considered the prediction of a point vortex system with observations and a data assimilation system. In their experiments they considered scenarios with two point vortices and four point vortices; they considered deterministic models and models forced by stochastic noise; they considered observations of vortex positions and fixed-point observations of the fluid velocity; and all experiments used an extended Kalman filter for the DA (see section 2.3.3). In their studies, they carved out the region of parameter space where prediction proceeded smoothly. Success, of course, depends on the frequency, density, and accuracy of available observations. However, they also were able to begin discerning between the value of so-called Lagrangian observations (i.e., vortex positions) and so-called Eulerian observations (i.e., fixed point “station” observations). In general, Lagrangian observations tend to strongly influence the update of their corresponding vortex, but only weakly influence the other vortices, whereas Eulerian observations influence all vortices in proportion to their nearness to the observing station. They note that as Eulerian observations are nonlinearly related to the state vector, they can actually have a deleterious effect on the resulting state estimate if further observations do not exist to remove ambiguities (e.g., other position observations).

In addition to demonstrating that prediction and filtering a point vortex system is indeed possible, Ide and Ghil draw two main conclusions about the use of Eulerian observations in the filtering process. First they note that a point vortex’s singularity is only an issue when a station observation of the fluid velocity is made within very close proximity of the point vortex (where the velocity goes to infinity). The second problem they note is when velocity observations are nearly zero. This can arise from observing a saddle point between vortices (i.e., the near cancelation of two large velocity contributions) or when the observing station is far from any of the vortices. They find that the impact of velocity observations on vortices is proportional to  $l^{-4}$ , where  $l$  is the distance separating the station from a vortex. To overcome the shortcoming of singularities, they repeated many of their experiments in a sequel paper examining Rankine vortices (Ide and Ghil 1997b). Rankine vortices act like point vortices outside of a pre-defined core region of solid body rotation, and hence eliminate the singularities of point vortices. To address the problem of small

velocity observations, they impose a minimum velocity cut-off criterion to represent a crude quality control scheme: if a velocity observation is less than the cut-off, then it is not used in the update. These changes largely addressed the problems from the first study.

More recently, Kuznetsov et al. (2003) considered the problem of assimilating time-integrated Lagrangian data as collected from drifters. They used a point vortex model with passive tracer particles, and they evaluated the tangent linear model of the point vortex equations and tracer equations so as to use an EKF. They found that tracer observations were enough to keep their estimate tracking truth's vortices within certain observational parameter ranges (e.g., accuracy, frequency, coverage). They also found a marked sensitivity to the initial tracer distribution. Their analysis shows the tracers to evolve according to underlying "Lagrangian flow structures" within the field. These Lagrangian flow structures, also called Lagrangian coherent structures, are invariants of the Lagrangian dynamics regardless of the frame of reference (e.g., Haller and Yuan 2000). This is an interesting topic that highlights the growing interest and utility in of mixing Lagrangian and Eulerian information in DA and forecasting. The frameworks established in this thesis could ultimately be applied to Lagrangian flow structure problems.

None of the studies have yet considered an ensemble approach to filtering a point vortex model, nor have they made the connection with an Eulerian state vector representation and the evolution of a Euler fluid.

### 4.2.3 Ensemble-based Kalman filtering

To confirm that an ensemble-based Kalman filter can perform well when used with a point vortex model, we begin with the simple four vortex configuration shown in figure 4-3. The intent is to run an Observation System Simulation Experiment (OSSE) under the perfect model assumption. This means that all evolution, both of the state designated "truth" and of the ensemble, is governed by the same model dynamics — truth and the ensemble differ only in their initial conditions. Truth is "observed" at a given time interval, termed the *observation time*. The observation types (i.e., position or velocity) are pre-specified. The only source of observation error is assumed to be instrument error, and hence, all observations are assumed to have uncorrelated errors of a pre-specified magnitude. Thus, observations are found by calculating the corresponding value from the truth state and adding a random sample from the prescribed observational error. Experiments are then run by varying the "observational network," including their coverage, frequency, and accuracies. To seed the initial ensemble, we first completely observe the true state vector (i.e., the initial positions shown in the left panel of figure 4-3) to form a base profile, and then we observe the

base profile  $N$  times, each resulting state being an ensemble member. This way, the initial ensemble's uncertainty is generally within observational uncertainty and consistent with truth, though truth is not forced to be at the center of the ensemble. An initial ensemble generated in this fashion will in general have zero initial statistically significant covariance between its state elements, but the ensemble integrations will quickly impart this.

The dynamics of a point vortex model can be nondimensionalized to yield a single control parameter. Rewriting the equations (4.4) and (4.5) as a single equation for the azimuthal velocity about a point vortex:

$$\frac{dv_\theta}{dt} = r \frac{d\theta}{dt} \sim \frac{\Gamma}{2\pi r}, \quad (4.6)$$

where strictly the righthand side is a summation. Substituting in nondimensional variables  $r \sim Lr'$ ,  $t \sim Tt'$ , and  $\Gamma \sim G\Gamma'$ , and dividing both sides through by  $\frac{L}{T}$ , yields:

$$r' \frac{d\theta}{dt'} \sim \frac{GT}{L^2} \frac{\Gamma'}{2\pi r'}. \quad (4.7)$$

This shows that the dynamics scales according to the parameter  $\frac{GT}{L^2}$ . Hence, if one changes the vortex strengths, the integration timestep, and the vortex separations by the same factor, one will get dynamically identical behavior over many scales. Our experiments use the same relative proportions of the separations in figure 4-3, however, the length scales are in terms of  $10^6$  m. We use four equal vortex strengths of  $1.0 \times 10^6$  m<sup>2</sup> s<sup>-1</sup> and a timestep of  $1.0 \times 10^5$  s. The model is integrated with a standard fourth order Runge-Kutta scheme.

As discussed above in section 4.2.1, the point vortex model is not a dissipative system, and as such has no attractor. Though neighboring trajectories are observed to diverge exponentially, the system is actually conservative meaning that there is always a compensating contraction in state space to counteract the observed divergence. We have coded the tangent linear model of the point vortex system, and calculated its Lyapunov exponents (see discussion in section 2.2.2). The system's eight exponents add to zero, again indicating its conservative nature. The largest positive exponent allows an estimate of the system's *error doubling time*, that is, the average time it takes an initially small error (meaning linear error dynamics and exponential error growth are valid) to double its magnitude. For our configuration and parameters, this calculation yields  $\tau_{\text{doub}} \approx 285\Delta t$ . Knowing the error doubling time allows a calibration and intercomparison of disparate dynamical systems. For instance, if one believed that the actual atmosphere's error doubling time were 2 days, then one could determine the number of model  $\Delta t$ 's that constitutes an "atmosphere day," about 140. Ehrendorfer (1997) and references within estimate the atmosphere's average error doubling time to be  $2.0 \pm 0.5$  days.

Since NWP runs a DA cycle every 6 to 12 hours, this suggests a relevant observation time to run with our OSSEs, namely four to eight cycles per error doubling time, or about every 20 to 40  $\Delta t$ 's. The expectation is that running DA this frequently should be enough to track the vortices well, given sufficient observational coverage and accuracies. We shall term this relative frequency of DA to a system's error doubling time "weakly nonlinear." It should be noted that one of the reasons features cause DA systems to fail is that their associated errors can violate the tangent linear hypothesis in a short duration of time (this will be demonstrated in the next chapter). Therefore, the error doubling time associated with features in the atmosphere ought to be much shorter than 2 days, and the relative frequency of DA to features' associated error doubling time probably falls within a regime of nonlinear error growth. This gives the expectation that we ought to be able to lengthen the observation time to place the system in a regime of nonlinear growth where the DA system's assumptions are violated, just as in the atmosphere.

Like Ide and Ghil (1997a), the observations considered include both position observations and station observations. As the point vortex model does not require discretizing space, to evaluate the observations at station locations, we could simply evaluate the field values at those points using equations (4.2) through (4.5). However, quite like Ide and Ghil found, we expect there to be issues with the point vortices' singularities, and since we are using an ensemble, we have many more opportunities for these issues to be raised. Therefore, when determining Eulerian field values at specified points, we assume a slightly different shape for the point vortices, namely Gaussian vortices. This has the same effect as when Ide and Ghil (1997b) chose to consider Rankine vortices, meaning that the vortices still appear to be point vortices to one another when their separations are large compared to their own length scale. We choose a Gaussian profile over the Rankine profile because we intend to insert them into a dynamical model, and a Gaussian profile is better behaved numerically. We construct our Gaussian vortices by assuming that their vorticity is contained within a two-dimensional Gaussian "hump" rather than a  $\delta$ -function. In terms of cylindrical coordinates about a vortex's center, we assume the form

$$\zeta(r) = A \exp(-Br^2). \quad (4.8)$$

We choose the constants  $A$  and  $B$  so that the vortex appears as a point vortex with associated azimuthal velocity  $v_\theta = \frac{\Gamma}{2\pi r}$  from a distance. In cylindrical coordinates,  $\zeta = \frac{1}{r} \frac{d}{dr} r \frac{d\psi}{dr}$

and  $v_\theta = \frac{d\psi}{dr}$ . Integrating this away from the center of the vortex, we obtain

$$v_\theta(r) = \frac{A}{2Br} (1 - \exp(-Br^2)), \quad (4.9)$$

from which we deduce that  $A = \frac{B\Gamma}{\pi}$ .  $B$  itself determines the Gaussian hump's length scale. We choose to define  $B$  to require that the Gaussian hump reaches one tenth its maximum value over a given distance,  $d$ , from its center. Hence,  $B = -\frac{\log(0.1)}{d^2}$ , where  $d$  is user-defined. Figure 4-4 shows a comparison of the azimuthal velocity profiles associated with three different vortex specifications: a point vortex, a Rankine vortex, and a Gaussian vortex. All vortices have  $\Gamma = 10$ , and the Rankine vortex and Gaussian vortex both assume a core size of  $d = 1$ . Note that outside the core, all profiles are essentially identical. An expression for the Gaussian vortex's  $\psi$  field is available in terms of the exponential integral,  $\int_r^\infty \frac{\exp(-x)}{x} dx$ , but we omit the exact form here. It is important to note that the Gaussian vortex approximation is only made when evaluating station observation values, the actual model dynamics being integrated are still equations (4.4) and (4.5).

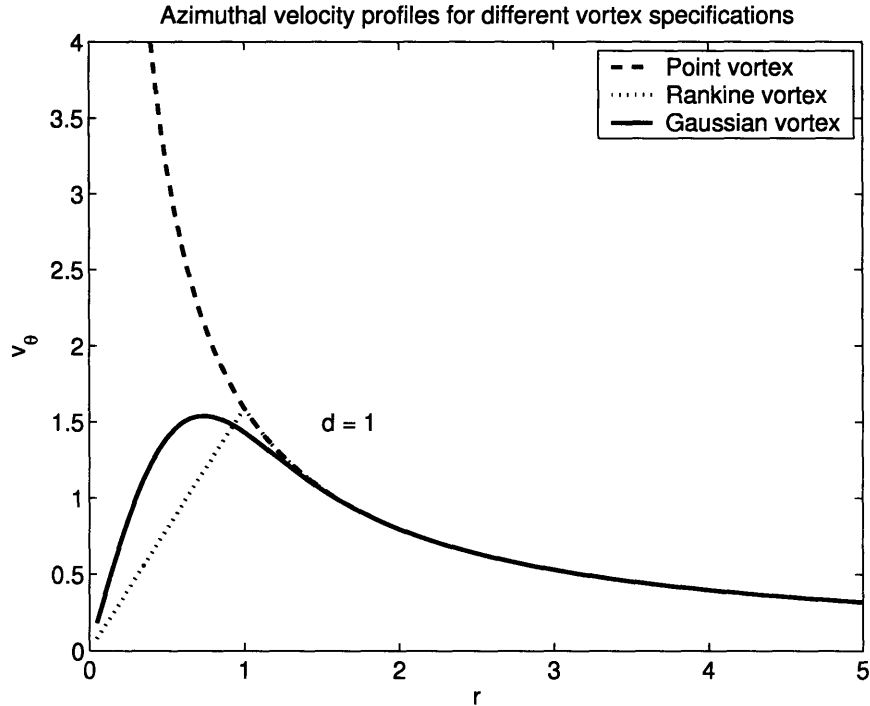


Figure 4-4: A comparison of three different azimuthal velocity profiles associated with different vortex specifications.

The experimental set-up we choose is a straightforward OSSE, initialized as described

above. We have run many different variations, and found that they show precisely what we expect, namely that a perturbed observations EnKF can adequately handle a nonintegrable configuration of a four point vortex system within a weakly nonlinear error growth regime. The success of course depends on the number and types of observations taken, the observation time, the observation accuracies, and the number of ensemble members used. We will discuss the typical parameters used in the OSSEs and present representative figures of their results in the next subsection where we compare the EnKF’s performance when applied to the Lagrangian state vectors and their corresponding Eulerian state vectors.

This success comes as no surprise. The way to view this result is as the consequence of treating the point vortex system, again in the spirit of Bjerknes, as “just another dynamical system” whose state vector happens to be  $\mathbf{x}$ . If the error PDF associated with  $\mathbf{x}$  is kept approximately Gaussian, and if  $\mathbf{x}$  is directly observed or adequately indirectly observed, again with approximately Gaussian error PDFs, then a Kalman filter based DA system will work. This may seem an obvious result, but it demonstrates the point that Lagrangian errors are easily corrected within a Lagrangian framework. Therefore, since one of the main problems associated with features springs from trying to update Lagrangian errors within an Eulerian framework, one approach to filtering in the face of features might be to transform the Eulerian dynamics into a problem with Lagrangian dynamics. Unfortunately, this is not such an easy feat. As mentioned above, one can try to approximate the equations for an Euler fluid using an aggregate of point vortices, but this does not render the DA problem any easier as one would need to draw correspondence between the many point vortices in one’s estimate and the available observations, or even more daunting, between the many vortices in each ensemble member. Indeed, it seems that only in the restrictive cases where one knows the underlying dynamics really are encapsulated by a simpler Lagrangian description, like the point vortex model, could one expect to successfully transform an Eulerian problem into a Lagrangian one. We conclude that transforming Eulerian problems into Lagrangian ones is not a general enough approach to consider further; however, in the case of the point vortex system where we know it will work, it provides a good testbed to see how much of an improvement could be expected were such a transformation technique available for the general case, that is, it provides a best case scenario.

#### 4.2.4 Comparing Eulerian and Lagrangian state vectors

In this section we compare the update of a point vortex state vector (i.e., the  $(x, y)$  positions of each vortex) with the update of its corresponding Eulerian state vector, the discretized vorticity field of the Euler fluid implied by our Gaussian vortex approximation. As the

success of the updates depends on the size assumed for the Gaussian vortices, we examine two different sizes in detail: relatively narrow vortices and relatively broad vortices. As explained in the previous section, the point vortex state vector ensemble performs well, and we find this is the case as long as they are spaced far enough apart relative to the Gaussian vortex length scale; conversely, we find that an ensemble-based Kalman filter performs better for the Eulerian state vector as the vortices become broader. The results in this section are kinematic in the sense that no dynamical models are used to cycle the obtained analyses though to the next observations time, though the point vortex model was used to evolve the states to the configurations shown. We will revisit the Eulerian representation of a point vortex model in a dynamic setting in chapter 6, and there we will show the effect of cycling on the obtained analyses.

### Experimental set-up

In this subsection, we use the results of the previous subsection, where the Lagrangian state vector of a point vortex model was filtered with an EnKF, as a benchmark to see how well an EnKF can update the same states represented over a gridded domain. Dynamically, the Lagrangian formulation evolves the vortex positions using their circulations as parameters of the model, whereas the Eulerian formulation explicitly evolves the continuous vorticity field represented over a grid. This means the vortex positions can, at best, be diagnosed from the Eulerian states, and this is how we proceed in the experiments in chapter 6. However, in the experiments presented in this subsection, we do not actually employ an Eulerian model to evolve the vorticity field, rather we use each point vortex state to evaluate the corresponding Eulerian field(s) at specified gridpoints using equations (4.2)–(4.5) whenever we are to perform DA. Foreshadowing a result to be made more explicit in the next chapter, here we find that the success when using an Eulerian state vector depends on the assumed size of the Gaussian vortex representations and on how dispersed the ensemble members are in position.

To compare the performance of an EnKF operating directly on the  $(x, y)$  positions of each point vortex against an EnKF operating on a gridpoint lattice of an Eulerian field associated with a given point vortex placement, one must first specify the dimensions and resolution of a grid. As stated above, the point vortex model is not bounded *per se*, and therefore defining a sensible gridpoint domain requires some foresight. Having explored the trajectory bounds of the above four point vortex set-up (see figure 4-3), we decide to use a square domain centered at the origin with a side length of  $4 \times 10^6$  m. Any resolution is possible (within computational bounds), but we chose to consider only 64 by

64, meaning  $\Delta x = \Delta y = 62.5$  km, or 128 by 128, meaning  $\Delta x = \Delta y = 31.25$  km. The comparison experiments examine changing the Gaussian vortex size parameter  $d$  (see the previous subsection). Thus for simplicity, we define  $d$  in terms of  $\Delta x$  or  $\Delta y$ . For instance, if one wants the Gaussian vortex profiles to decrease from their maximum values to one tenth those values within  $N_d$  gridpoints, then one uses  $B = -\frac{\log(0.1)}{(N_d \Delta x)^2}$  in equation (4.8).

Since the point vortex model derives from the equations for an Euler fluid, whose evolution is completely specified by either its vorticity field or its streamfunction field, its proper Eulerian description should yield a univariate state vector, either  $\zeta$  or  $\psi$  — balance between  $u$  and  $v$  is automatically assumed. If one were interested in seeing how far out of balance DA brings the analysis, then following Ravela et al. (personal communication, 2004), one could assume a multivariate, “geostrophic” state vector of  $u$ ,  $v$ , and  $p$  (where  $p$  is pressure), and then diagnose the degree of imbalance, but this is an unnecessary complication for the point we are trying to make here. Using  $\zeta$  or  $\psi$  for the state vector is equivalent; we choose to use  $\zeta$  as our state vector out of convenience (e.g., the conversion from point vortex positions to  $\zeta$  is much easier for a Gaussian vortex than to  $\psi$ ). If one did not use the Gaussian vortex approximation, then the  $\zeta$  field would contain only  $\delta$ -functions, leaving essentially the entire field with zero vorticity. When one makes the Gaussian vortex approximation with small parameter  $d$ , then one still expects most of the  $\zeta$  field to be nearly zero. This means that direct observations of the state will give very little information to the updating process unless the observing station is very close to one of the vortices. This is okay because observations of  $\zeta$  are not physical —

unless one were to take an approach similar to Li et al. (1998) and assert that passive tracers were decent proxies for vorticity data (see section 2.4.3). Instead we choose to indirectly observe the state through  $u$  and  $v$ .

As in the previous subsection, we again consider station observations of velocity and position observations of vortex position. The observation operator that relates a gridpoint representation of  $\zeta$  to vortex positions is quite nonlinear. If one deems a vortex’s position to be defined by the location of the maximum in vorticity, then the function,  $\mathcal{F}(\zeta)$ , whose output is the location(s) of the maximum value in vorticity, is essentially evaluating the function  $\arg \max(\zeta)$  (also sometimes called *max loc*). Were one to try to linearize this function about the state, one would find a matrix whose elements were mostly 0 punctuated by a few  $\pm\infty$ ’s. This means that traditional Kalman filter methods that rely upon evaluating the Jacobian of the observation operator,  $\mathbf{H}$ , will not be able to use position observations to update an Eulerian state vector. There is a well known way around this problem for filtering exercises which we will explain in detail in chapter 6, namely using an augmented

state vector.<sup>2</sup> Ensemble methods need not use an augmented state vector because the required covariance products involving  $\mathbf{H}$  within the expression for the Kalman gain matrix can equivalently be estimated directly from the nonlinearly observed ensemble (as described in section 2.3.3).

Once the DA has been performed, there are two ways to compare the analysis ensembles: within the Lagrangian framework and within the Eulerian framework. The Lagrangian framework compares vortex positions, whereas the Eulerian framework compares vorticity fields. To compare estimates within the Lagrangian framework, one must try to deduce the locations of the updated vortices in the gridpoint-based ensemble to compare to the point vortex analysis. This can become a subjective task as the Eulerian update deforms the vorticity field into one with multiple local maxima. The Eulerian framework requires that the point vortex ensemble be converted to a gridpoint based vorticity field, the same procedure undertaken to form the Eulerian ensemble to begin with. In practice, comparison within the Eulerian framework is both easier (i.e., less subjective) and closer to what fluid dynamicists would likely be interested in. We include an example of both comparisons below.

### Narrow vortices, $d = 3\Delta x$

#### *Prior ensembles*

Here we considered examples with relatively narrow vortices. Because the point vortex model and its gridpoint based representation occupy two spatial dimensions, visualizing ensembles of their output can be tricky. One well-known technique we employ is termed a “spaghetti diagram,” and it consists of plotting the same single contour from each ensemble member conterminously on a map. Though this does not present the entire field, one is able to get a sense of the ensemble dispersion *in position* on one figure. In the case of a point vortex model, position dispersion is all that matters, so the spaghetti diagram is ideally suited to our needs. Figure 4-5 shows one such example of this. The left panel shows a short-term forecast ensemble with  $N = 50$  members from the point vortex model. The axes of this panel and the others to follow are all labeled in units of  $10^6$  m, just for clarity. Each of the four vortices is plotted in a different color, therefore allowing quick visual correspondence of vortices among the ensemble members. The ensemble was spun-up from its initial state through many successive DA cycles operating within the weakly nonlinear parameter range discussed in the previous subsection. At each observation time, only the

---

<sup>2</sup>Augmented state vectors will still not allow adjoint based state estimation to proceed, because the adjoint of the function  $\mathcal{F}$  is filled with either 0’s or  $\infty$ ’s.

blue vortex's position was observed, with an accuracy of  $\sigma_{\text{pos}} = 1.5\Delta x$ , where  $\Delta x$  is set by the 64 by 64 grid resolution (i.e.,  $\Delta x = 62.5$  km). There were also ten station observations made of the fluid velocity at randomly, pre-selected gridpoint locations with an accuracy of  $\sigma_{\text{vel}} = \frac{1}{10}V_{\text{max}}$ , where  $V_{\text{max}}$  is the maximum azimuthal velocity associated with a single vortex.  $V_{\text{max}}$  is set by the parameter  $d$  in the Gaussian vortex profile. In this case we have chosen  $N_d = 3$ , meaning that  $V_{\text{max}} \approx 0.822 \text{ m s}^{-1}$ , leading to an observational uncertainty of  $\sigma_{\text{vel}} = 0.0822 \text{ m s}^{-1}$ .<sup>3</sup> The ensemble estimate was able to successfully track truth's evolution throughout the initial spin-up time. After the spin-up period, the observation time was lengthened to  $80\Delta t$ . This allows the ensemble dispersion to become larger. Figure 4-5 shows a snapshot of the forecast ensemble after five DA cycles past the spin-up period. The points show the dispersion, and we see that the black and blue vortices have spread in the  $y$  direction comparable to the domain length. To be sure, this is quite a difficult case to correct given the few observations available, and it will highlight many of the issues presented by state estimation in the face of features eluded to above, especially in the heuristic example shown in figure 4-1.

The points in the left panel of figure 4-5 are the individual ensemble members from the point vortex model. We denote this ensemble as  $\mathbf{Z}_L^f$  (subscript "L" for Lagrangian), a matrix whose columns are the  $N$  different ensemble members, each with 8 state elements. Truth's positions are shown by the thick yellow circles. Truth's positions are not color-coded because their correspondence to the other vortices is clear. Each colored grouping within the ensemble admits its own mean vortex position, and these are shown by the like-colored diamonds. The mean positions are at the center of like-colored ellipses. The ellipses show the covariance of uncertainty (with one standard deviation per ellipse axis) between the  $x$  and  $y$  positions *for each vortex*. There is strong covariance between the positions of different vortices as well, but this information is difficult to display in a two-dimensional figure. To gain a sense of the covariance, note that truth's positions are all rotated counterclockwise with respect to the ensemble mean positions; the individual ensemble members are similarly related to the mean. It is clear from looking at either the black or blue vortices that their position distributions are not sufficiently described by only their means and covariances. For instance, the black vortices carve out a slightly crescent shaped distribution. This indicates that if an ensemble member has a  $y$  value for its black vortex that it is either much greater or much less than the mean  $y$  value, then it is likely its corresponding  $x$  value will be greater than the mean  $x$  value. This correlated structure requires more than two moments of PDF

---

<sup>3</sup>Differentiating equation (4.9) and solving for its maximum value is not easy in general. We have found that to a very good approximation, the maximum velocity is  $V_{\text{max}} \approx (\Gamma B^{\frac{1}{2}})/9.8456$ .

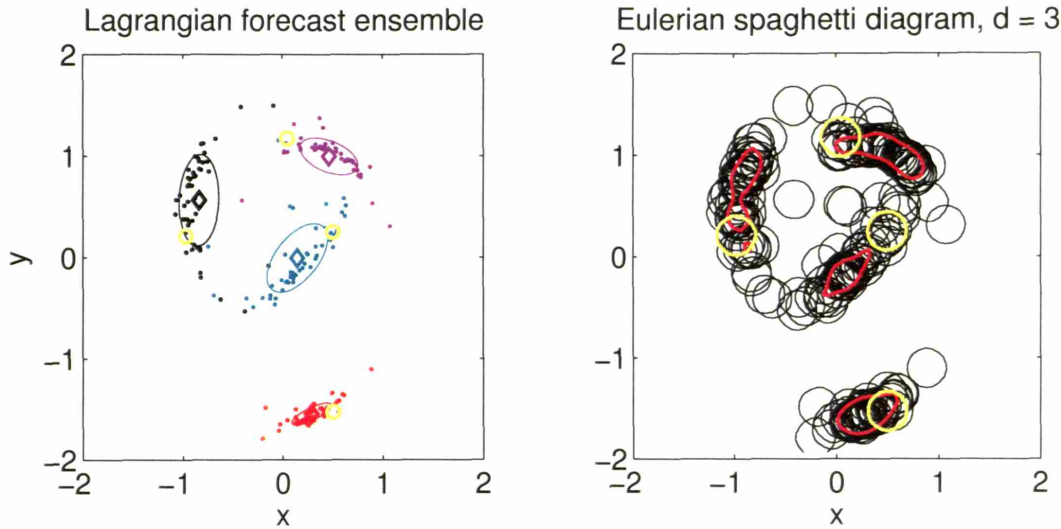


Figure 4-5: The left panel shows a widely dispersed short term forecast ensemble forecast from the point vortex model. Corresponding vortices are like-colored. Truth's positions are shown by the thick yellow circles. The ensemble mean positions are denoted by diamonds, and the covariance for each vortex between its  $x$  and  $y$  positions are denoted by the ellipses. The right panel shows a spaghetti plot constructed from the gridpoint representations of the ensemble members in the left panel. The contour value show corresponds to one tenth of the vortices' maximum value — hence, the circular contours have a radius of  $3\Delta x$ . The yellow contours correspond to truth, and the red contours are from the ensemble mean.

to represent it. Also note that the black vortex's mean is almost outside of its ensemble grouping. It is interesting to note that these non-Gaussian vortex distributions each look very similar to the state space snapshots of the Ikeda system examined in the last chapter (*cf.*, figure 3-8). Hence, we suspect that the ensemble updates could be bad if there are not accurate enough observations or strong enough covariances between vortices to guide the estimation process. That the vortices have non-Gaussian distributions themselves shows the entire ensemble has a non-Gaussian distribution in its state space; however, the reverse is not true, Gaussian individual position distributions are not enough to guarantee that the vortices do not vary non-normally between one another.

We denote the process of transforming from point vortex model space to gridpoint model space by the operator  $\mathcal{P}$ , of which  $d$  is a parameter, and we denote the ensemble in gridpoint space as  $\mathbf{Z}_E$  (for Eulerian). Hence,  $\mathbf{Z}_E^f = \mathcal{P}[\mathbf{Z}_L^f; d]$ . The right panel of figure 4-5 shows a spaghetti diagram for the gridpoint representations of the Gaussian vortices ( $\mathbf{Z}_E^f$  with  $d = N_d\Delta x$ ) associated with each of the point vortex locations in the left panel, that is, the left panel shows the Lagrangian state vector ensemble and the right panel shows

the corresponding Eulerian state vector ensemble. The contour chosen in the spaghetti diagram has the value of one tenth the vortices' maximum vorticity value (given by the constant  $A$ ). Each black circular contour corresponds to a point in the left panel. The thick yellow contours show the same contour value for truth, and the thick red contours show the same contour value of the ensemble mean. Note that the ensemble mean contours are not circular. The spaghetti diagram is perhaps more imposing than the left panel because the color-coded vortex correspondence has been lost. The ordering and labeling of vortices is explicit in the point vortex model, but the gridpoint representation shows only the two-dimensional distribution of  $\zeta$ , thus no significance is explicitly given to features. Note also that the Eulerian state vectors each need 4096 elements at this given resolution to represent what the point vortex model encases in only 8 state elements (i.e.,  $2^{12}$  versus  $2^3$  elements).

### *Observations*

Figure 4-6 shows the observations taken for this DA cycle. Both panels show the ten station velocities taken and the one position observation. The left panel shows these observations in relation to the point vortex model. As in figure 4-5, truth's positions are shown by the thick yellow circles, and the ensemble mean positions are shown by the colored diamonds. The vector arrows show the velocity observations with magnitudes relative to one another's strengths. The velocities are generated by assuming that truth really has a Gaussian vortex profile. The brown asterisk shows the position observation of the blue vortex. The right panel shows the same observations in relation to the gridpoint representations. Instead of plotting a spaghetti diagram, here we show the ensemble mean (thick black contours) superposed on the ensemble spread (shaded regions). The ensemble spread is simply the standard deviation of each gridpoint value as represented through the ensemble. Spread below a certain value has been omitted. The contour interval for the ensemble mean is in increments of one tenth of the vortices' maximum value. Truth is represented by the thin yellow contours with the same contour interval (i.e., truth has all ten contours shown). The fact that the ensemble mean only has at most four concentric contours shows how much smearing has been accomplished by averaging ensemble members with wide position dispersion.

### *Analysis ensembles*

Figure 4-7 shows the EnKF analyses of the two initial ensembles shown in figure 4-5 using the observations shown in figure 4-6. The left panel shows the Lagrangian state vector analysis,  $\mathbf{Z}_L^a$ , as considered in the previous subsection (section 4.2.3, though figures are not

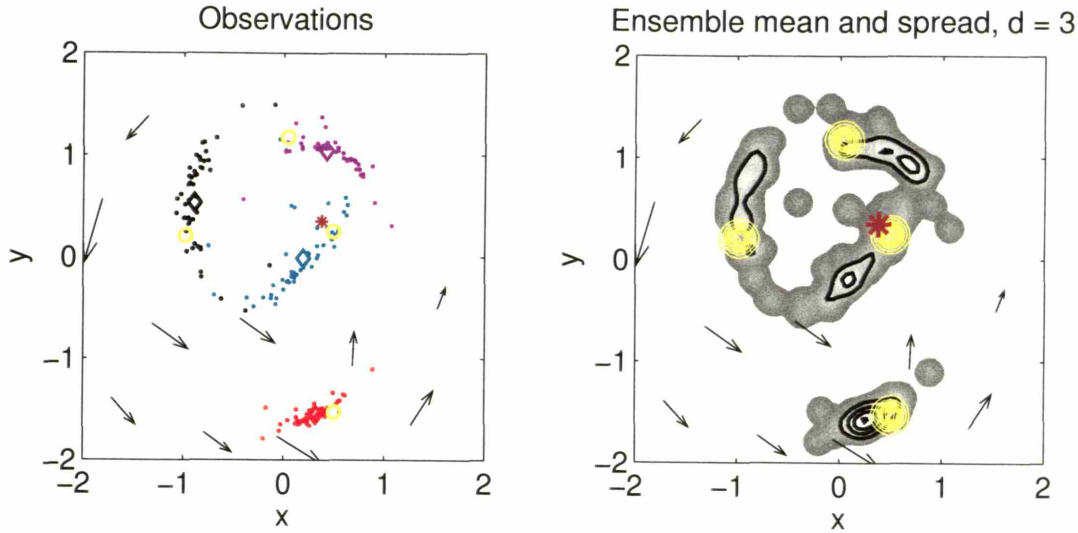


Figure 4-6: The left panel shows the point vortex forecast ensemble, its mean, and truth as seen before in figure 4-5. The black vector arrows show the ten station observations of fluid velocity available for DA taken from assuming that truth has a Gaussian vortex profile. The brown asterisk shows the one position observation taken of the blue vortex. The right panel shows the same observations as in the left panel. The shaded region shows the ensemble spread, and the thick black contours show the ensemble mean. The thin yellow contours show truth's vorticity profile.

shown there so that the figures here are not redundant). Note that much of the original non-Gaussianity has been eliminated. It is interesting to see just how effective the EnKF was in this example, especially given the non-Gaussian prior distribution and the limited observation information. Further experimentation shows that the position observation is the most valuable for informing the ensemble update. The covariance that exists between vortices allows a single observation to affect the analysis position of all four vortices. By random chance, there are not that many observing stations in the upper half of the domain, and their effect on the black, magenta, and blue vortices is weak. It turns out the non-Gaussianity is not too much of an issue in this case because the prior distribution is rather uncertain compared to the position observation and the strong covariance between vortex positions. Truth's positions are therefore accordingly consistent with the ensemble's estimate thereof.

The right panel of figure 4-7 shows the spaghetti diagram made from  $\mathbf{Z}_E^a$ .  $\mathbf{Z}_E^a$  is the analysis ensemble from applying an EnKF directly to  $\mathbf{Z}_E^f$  (i.e., directly updating the grid-point representations). The spaghetti contour value is the same as was used for the prior ensemble in the right panel of figure 4-5, and truth's contours are included as a reference for

comparing the black contours between spaghetti diagrams. We see that the ensemble mean contours of  $\mathbf{Z}_E^a$  are in better agreement with truth than are those of  $\mathbf{Z}_E^f$ . However, there is still quite a lot of spread in the black contours, about as much as in figure 4-5, only now the contours are of variable size. Essentially every initial contour in figure 4-5 still has a signal in this figure. This is partly because the contour we chose to consider for the spaghetti diagram, one tenth of the maximum vorticity value, has such a low value. We chose this because it also allows visualization of the ensemble mean, whose maximum vorticity value is much less than any of the individual members due to smearing effects from averaging (*cf.*, figure 4-1). Were one to choose a contour value of, say, half the vorticity maximum, then there would be more of an obvious difference between this figure and figure 4-5, but we would not be able to see the ensemble mean. Regardless though, it is apparent that the updated Eulerian ensemble has significant amplitude in places where there are no vortices shown in the Lagrangian ensemble in the left panel.

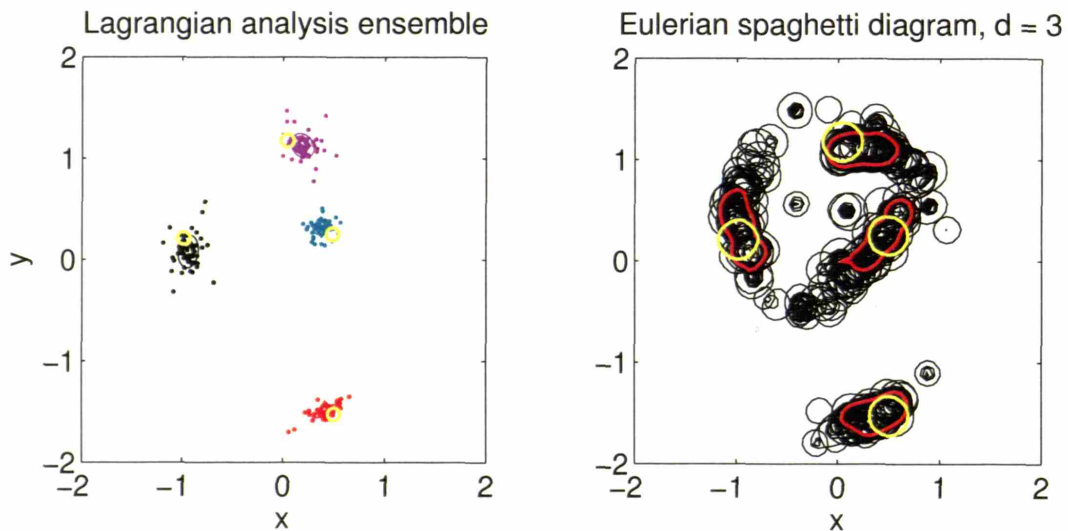


Figure 4-7: The left panel shows the point vortex analysis ensemble, its mean, and truth as updated from that shown in the right panel of figure 4-5. The right panel shows the spaghetti plot of the analysis ensemble of gridpoint representations. The contour value considered is one tenth truth’s vorticity maximum. The black contours are from the individual members, the yellow contours are from truth, and the red contours are from the ensemble mean.

#### *Cross-validation analysis ensembles*

By way of comparing the two analyses shown in figure 4-7, we generate the two cross-validation analyses referred to in the experimental set-up section above, that is, comparison

within the Lagrangian framework and comparison within the Eulerian framework. These cross-validation analyses are shown in figure 4-8. The left panel shows the analysis for comparison within the Lagrangian framework: it is our best effort at diagnosing the vortex positions from the Eulerian analysis ensemble (diagnosing position from  $\mathbf{Z}_E^a$ ). This panel should be compared to the proper Lagrangian analysis shown in the left panel of figure 4-7. It turns out that the operator  $\mathcal{P}$ , which transforms from Lagrangian state vector space to Eulerian state vector space, does not have a well defined inverse, so generating the Lagrangian cross validation analysis is not strictly objective. But by manually examining the individual ensemble members of  $\mathbf{Z}_E^a$ , we are able to present this subjective version of  $\mathcal{P}^{-1}[\mathbf{Z}_E^a]$ . Note that the analysis positions in this plot more closely resemble the initial Lagrangian ensemble,  $\mathbf{Z}_L^f$  (see figure 4-5), than the analysis Lagrangian ensemble,  $\mathbf{Z}_L^a$  (see figure 4-7). This is hinted at in the spaghetti diagram from which  $\mathcal{P}^{-1}[\mathbf{Z}_E^a]$  has been estimate (right panel of figure 4-7), but it is made very clear here.

The right panel of 4-8 is for the comparison within the Eulerian framework: it is a spaghetti diagram compiled from  $\mathcal{P}[\mathbf{Z}_L^a; d]$ , the Eulerian representation of the Lagrangian state vector analysis. This panel should be compared to the proper Eulerian analysis shown in the right panel of figure 4-7. The spaghetti contours in this figure are clearly much less dispersed than in the corresponding plot of  $\mathbf{Z}_E^a$  in the right panel of figure 4-7.

### *Individual ensemble updates*

As was previously mentioned, the updated individual ensemble members can readily lose their coherent vortex representation in the analysis process. We first present examples of this, and then seek to explain it afterwards. Figure 4-9 shows the prior and analysis states of three particular ensemble members. The left column of panels shows the Eulerian state vector forecast states,  $\mathbf{Z}_E^f$ , and the right panels show the corresponding analysis states,  $\mathbf{Z}_E^a$ , from direct application of an EnKF to the prior Eulerian state vector ensemble. The red contours show the vorticity in increments of  $\frac{1}{10}\zeta_{\max}$ , the yellow contours show only the  $\frac{1}{10}\zeta_{\max}$  contour of truth's vorticity, and the black circles show the corresponding values of the Lagrangian state vector forecast ensemble,  $\mathbf{Z}_L^f$  (i.e., the point vortex positions). In the right column of panels, the blue contours denote negative values of vorticity with the same contour interval as the red contours (the zero contour has been omitted), and the black dots denote the corresponding analysis positions of the point vortex model states,  $\mathbf{Z}_L^a$ . The three ensemble member updates were chosen deliberately to showcase the range of behavior. The top row denotes our subjective judgment of the “average update” as determined by careful inspection of all 50 ensemble members (and many other members in other experiments not

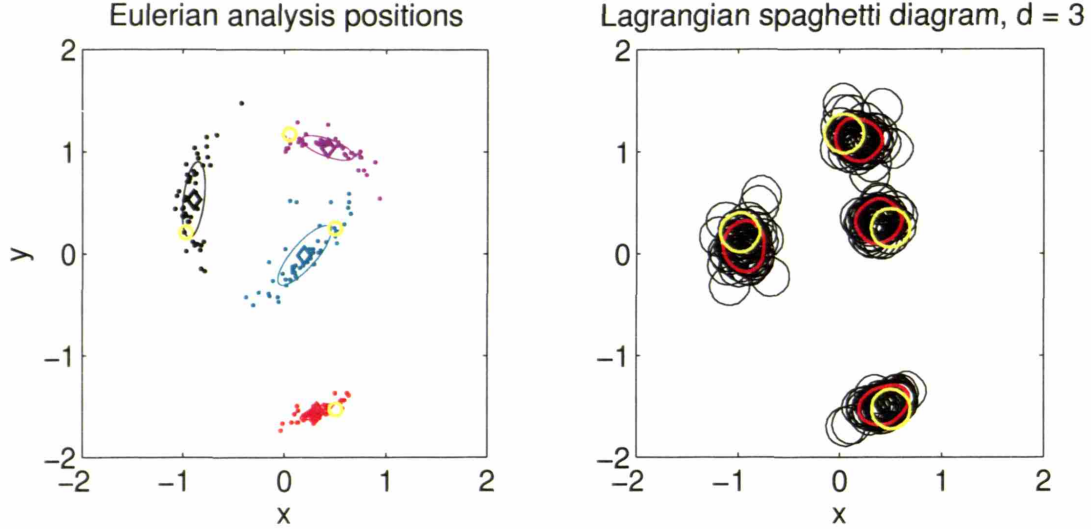


Figure 4-8: The left panel shows an estimate of  $\mathcal{P}^{-1}[\mathbf{Z}_E^a]$ , the diagnosed vortex positions from analysis Eulerian ensemble (as opposed to  $\mathbf{Z}_L^a$ ). Also shown is their mean and truth. The same colors and plotting conventions hold as in the previous three figures. The right panel shows the corresponding spaghetti diagram from  $\mathcal{P}[\mathbf{Z}_L^a; d]$  (as opposed to  $\mathbf{Z}_E^a$ ). Again, the contour value considered is one tenth truth’s vorticity maximum. The black contours are from the individual members, the yellow contours are from truth, and the red contours are from the ensemble mean.

presented in this thesis). The middle row of panels denotes what we term a “bad update,” and the bottom row of panels shows a “good update.” Before characterizing the analyses, we note that good, bad, and average here generally relate to how accurate the prior ensemble member is. In the case of the bottom row, the “good” case, the left panel shows there to be very good agreement with the vortex placement, whereas the middle row of panels, the “bad” case, shows the initial member is quite a bit in error.

Though glancing at the prior ensemble distribution shown in figure 4-5 leads one to suspect that there are many ensemble members in grave error, it turns out that because of the Gaussian vortex profile associated with  $d = 3\Delta x$ , a fair number of ensemble members are generally close to truth’s positions. The top left panel of figure 4-9 shows such a case. The estimate’s vortices are rotated slightly clockwise about the origin relative to truth’s positions (at least the top three vortices are). The right panel shows that the Lagrangian analysis positions,  $\mathbf{Z}_L^a$  (the black circles), have been rotated counterclockwise back toward truth’s positions, indicating the EnKF’s success in recognizing how the forecast was in error. It seems that the Eulerian analysis positions, as seen in  $\mathbf{Z}_E^a$  (the contours), have also rotated slightly counterclockwise relative to the prior positions, however, the corrective

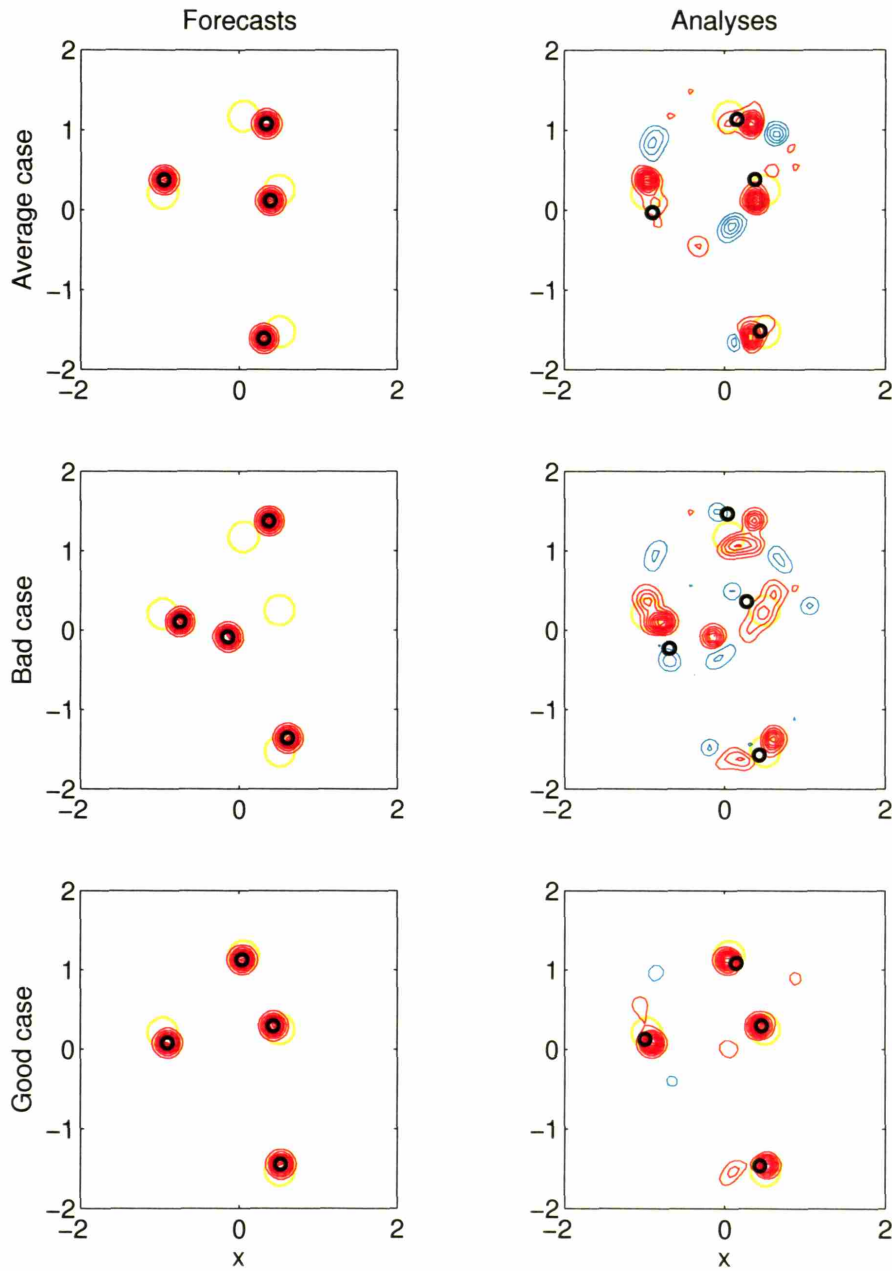


Figure 4-9: The left column of panels shows three specific prior ensemble members. The red contours are positive vorticity contours with contour interval  $\frac{1}{10}\zeta_{\max}$ . The thick yellow contours correspond to one contour of truth, and the black circles denote the corresponding values of  $\mathbf{Z}_L^f$ . The right column of panels shows the updates of the three ensemble members shown in the left column. The blue contours show negative vorticity values, and the black circles now denote the corresponding values of  $\mathbf{Z}_L^a$ . The top row of panels shows a typical ensemble member update, the middle row shows a particularly bad update, and the bottom row shows a rather good update.

process has caused significant “pools” of negative vorticity in the shifted vortices’ wakes (i.e., the emergence of blue contours immediately clockwise of the vortices being rotated counterclockwise). It seems this feature is fairly typical of trying to coherently shift vortices in one direction or another. This effect will be studied further and explained in the next chapter. Note that the bottom row does not suffer too notably from this effect because the prior member does not need very much shifting.

The middle row of figure 4-9 shows a bad update. However, it is difficult to present just one bad update because there are many ways an analysis can be bad, depending on what is most important to the user — the example shown here exhibits several of the less desirable attributes. One possibly bad attribute, shared also by the average case, is the emergence of regions with negative vorticity. None of the prior ensemble members have negative vorticity, yet the statistics used by the EnKF are not able to “recognize” this nor enforce it. Another bad attribute is that the four strongest vorticity maxima do not constitute a good estimate of truth — indeed, the four vorticity maxima in the prior are hardly altered in the analysis. This is the case as well for the good and average cases, however, in those cases the prior estimates had relatively small initial errors. In the bad case, the analysis process has seemingly generated multiple vorticity maxima, or alternatively, split the four original vortices into “double core” vortices. One can imagine that if this state were used as the initial condition for an ensuing forecast, as ensemble filters would dictate, then its evolution would be quite different from truth’s (we shall see an example of this in chapter 6).

### *Discussion*

Many of the bad aspects noted above are related to a commonly used measure of misfit, namely root mean square (rms) error. Rms error is an  $\ell^2$  norm concept (the Euclidean distance between two points in the corresponding state space). When one compares the analysis ensemble members to truth, one finds that the members with the highest rms error are those that have four intact, coherent vortices in the wrong locations (no such members are shown). Based on physical plausibility, having four coherent vortices is good, but based on the rms error measure of accuracy, having the vortices in the wrong place is bad. Conversely, based on physical plausibility, analysis members like the one shown in the right middle panel are undesirable, but they actually have lower rms errors because the smearing that has occurred tends to shorten the distance between truth and the ensemble member. This smearing is a form of “hedging,” where DA expects the member to be in error and is playing for minimal error rather than striving for a physically plausible, error free estimate. This smearing leads to other negative effects. One is that vorticity maxima are weakened,

and another is that the vortices' shapes can be altered from being round. Note that this occurs even though every ensemble member has vortices of the same size and the same strength: feature position discrepancy within a gridpoint based estimate does not allow access to this information. Also note that we have raised the issue of physical plausibility of the individual members, to enforce this consideration on the estimation process is to impose a further constraint on what we consider a good state estimate. We will return to this point in the next section, and we will discuss it at length in the subsequent chapters.

Understanding why this occurs and accounting for it is the main focus of the next two chapters, so we will not go into too much detail on it now. However, we will note that contributing to its occurrence is the fact that the gridpoint based errors generated by position discrepancies can readily become non-Gaussian. Also, as emphasized by writing the EnKF update in terms of its mean and its ensemble perturbations (see equations (2.42) and (2.43) or equations (3.1) and (3.2)), the filter is trying to find an analysis ensemble mean that approximates the minimum variance estimate of a presumed Gaussian PDF's mean. The individual ensemble perturbations are updated to ensure that the analysis mean has approximately the correct uncertainty associated with it. To reconcile this, one can either argue that there is no explicit enforcement of the members being physically viable, or one can argue that if the assumptions of linearizable dynamics and Gaussian PDFs were truly met, then all states would be physically viable (from simple linear superposition of ensemble perturbations and the mean). Clearly, these problems arise from the fact that we are applying linear/Gaussian machinery to a nonlinear/non-Gaussian problem.

Hence, the EnKF seeks to update the ensemble mean at the expense of the ensemble perturbations. Unfortunately, when the involved PDFs are non-Gaussian, the mean itself is a "bad" estimate (though possibly not in an rms error sense). To summarize this point and the results of the above experiments, figure 4-10 shows a comparison of the ensemble means (as opposed to the spaghetti diagrams shown above) of the Eulerian state vector updates and the gridpoint representations of the Lagrangian state vector updates (i.e.,  $\mathbf{Z}_E^a$  versus  $\mathcal{P}[\mathbf{Z}_L^a; d]$ ). Note that the Eulerian state vector analysis mean in the left panel is quite smoothed out with elongated features and does not resemble truth too closely. However, the Lagrangian state vector analysis mean in the right panel has kept relatively tight, compact features in about the correct positions. The right panel is both physically plausible and accurate. If one were evolving  $\mathbf{Z}_E$  instead of finding it at each time through the  $\mathcal{P}$  operator, then the analysis in the right panel shows how well one could do by transforming the Eulerian problem at hand to a Lagrangian problem. We will return to this problem in chapter 6 where we will actually evolve  $\mathbf{Z}_E$ .

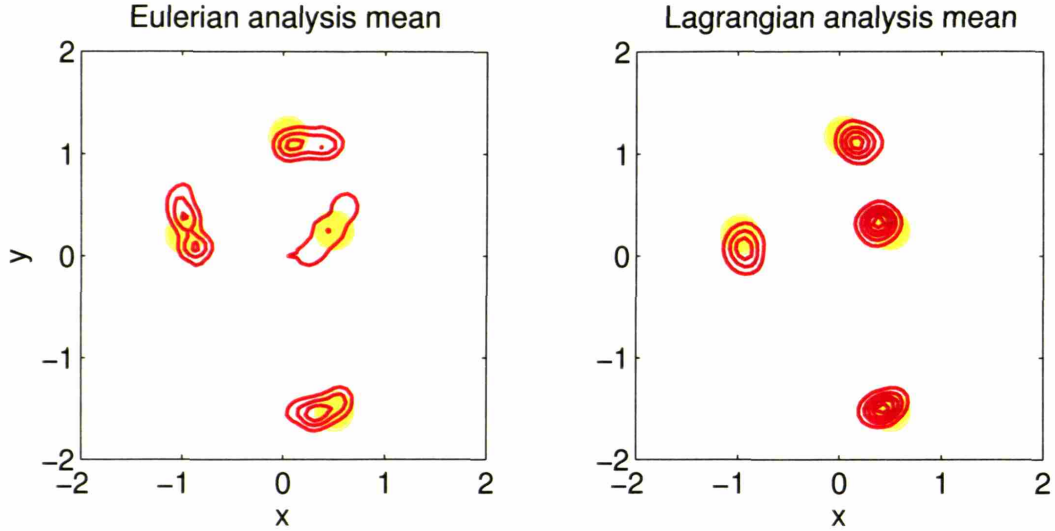


Figure 4-10: The left panel shows the ensemble mean of  $\mathbf{Z}_E^a$ , and the right panel shows the ensemble mean of  $\mathcal{P}[\mathbf{Z}_L^a; d]$ . The red contours are in increments of  $\frac{1}{10}\zeta_{\max}$ , and the yellow contours correspond to truth.

### The sensitivity to $d$

Before leaving this example, it is instructive to point out that part of what makes the gridpoint based updating difficult is the compact coherency of the features. Intuitively, it seems that by introducing another pertinent length scale into the problem through the parameter  $d$ , that its size might affect the results. It turns out that it does, and it is  $d$ 's size in relation to the ensemble position spread that matters most. To see this, figure 4-11 shows the prior Eulerian state vector ensemble mean (as calculated from  $\mathcal{P}[\mathbf{Z}_L^f; d]$ ) compared to truth for  $d = N_d \Delta x$ , where  $N_d = 1, 9$ . There are no updates shown — all panels show different Eulerian representations (each for a different value of  $d$ ) of the same initial Lagrangian ensemble (the one shown in the left panel of figure 4-5). The contour intervals are not comparable between panels because of the dependence of  $\zeta_{\max}$  on  $d$ . Also, the contour intervals are not comparable between the black contours of the ensemble mean and the red contours of truth because we wanted to show ten contours of each field to get the main structure across rather than emphasize the weakening and smearing that has occurred relative to truth (as we did above). The main aspect to note is the relative areal coverage of truth over the ensemble mean. As  $d$  increases and the vortices get larger, their distinct featuredness diminishes. Note, however, that as  $d$  gets really large, the Gaussian vortex approximation begins to break down, meaning that the vortices are so close to one another that they no longer appear to each other as point vortices. By this continuum of

$d$ , we see that point vortices really are the ultimate compact features since their  $d \rightarrow 0$ .

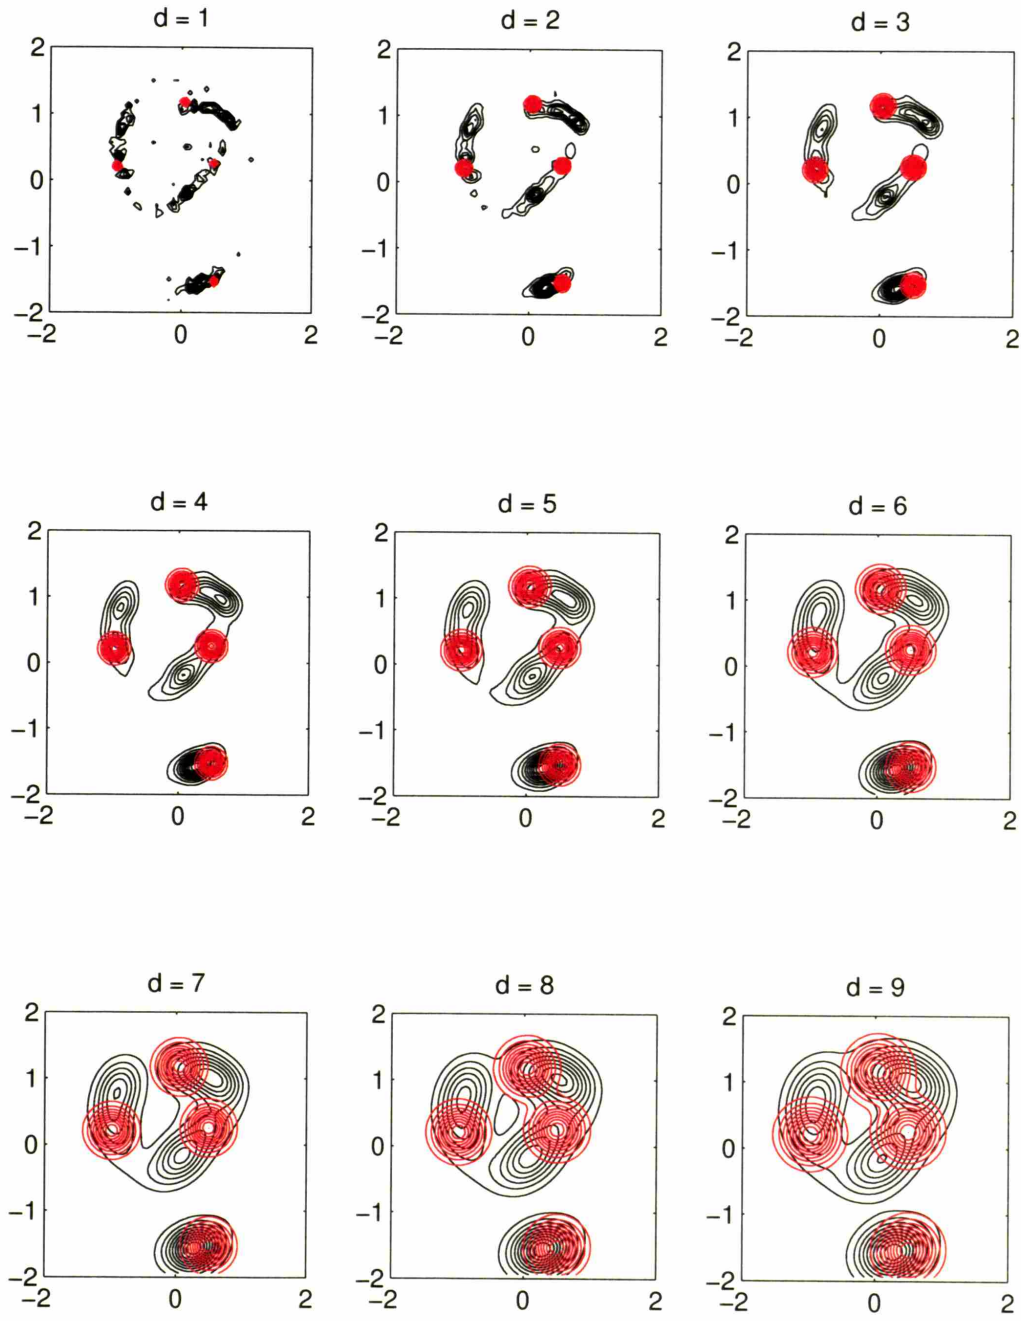


Figure 4-11: These nine panels show how  $\mathbf{Z}_E^f$  depends on the parameter  $d$  by way of the operator  $\mathcal{P}$ . Each panel shows the ensemble mean and truth for a given value of  $d$  (strictly the labeled values are for  $N_d$ , i.e., the number of  $\Delta x$ 's in the Gaussian vortices' cores). Truth is shown by the red contours, and the ensemble mean by the black. The contour intervals here are not comparable between plots, nor are the black and red contours comparable within a single plot.

### Broad vortices, $d = 6\Delta x$

Figures 4-5 through 4-10 are all for the case where  $d = 3\Delta x$ . This corresponds to the upper right panel in figure 4-11. To see how the updating improves as the featuredness diminishes, we present the salient plots for the case when  $d = 6\Delta x$ . Figure 4-12 shows six plots that address the ensemble distribution and update as a whole. Changing  $d$  will affect the gridpoint representations of the Lagrangian state vectors. This includes truth's representation, which means that the station velocity observations will change, and thus the Lagrangian state vector analysis will change. Hence, figure 4-12 includes plots of: the Eulerian state vector forecast ensemble ( $\mathbf{Z}_E^f$ , top left), the Eulerian state vector analysis ensemble ( $\mathbf{Z}_E^a$ , top right), the Eulerian forecast ensemble mean and the observations used ( $\bar{\mathbf{Z}}_E^a$ , middle left), the Eulerian analysis ensemble mean ( $\bar{\mathbf{Z}}_E^a$ , middle right), and for comparison, the Lagrangian state vector analysis ensemble for the new observations ( $\mathbf{Z}_L^a$ , bottom left), and its corresponding gridpoint representation ( $\mathcal{P}[\mathbf{Z}_L^a; d]$ , bottom right).

The top left panel shows a spaghetti diagram for the Eulerian state vector forecast ensemble ( $\mathbf{Z}_E^f$ ), but this time the contour considered is  $\frac{1}{2}\zeta_{\max}$  (as opposed to  $\frac{1}{10}\zeta_{\max}$  as used in figure 4-5). The other two spaghetti diagrams in figure 4-12 (the upper and the lower right panels) use this new contour value as well. The middle left panel shows the forecast ensemble mean, the ensemble spread (shaded), the available velocity observations, and the position observation. Also included are yellow contours showing the  $\frac{1}{10}\zeta_{\max}$  contour of truth to give an idea of the position and size of its vortices. Otherwise, all plotting conventions are the same as in the above figures where we consider  $d = 3\Delta x$ . Comparing the Eulerian state vector analysis in the upper right panel to the gridpoint representation of the Lagrangian state vector analysis in the lower right panel, we see their spaghetti diagrams look more similar than they did with  $d = 3\Delta x$  (*cf.*, figures 4-5 and 4-8). Also, the analysis ensemble mean shown in the middle right panel is tighter and more similar to truth.

To gauge whether the individual ensemble members also benefit from the reduced featuredness, figure 4-13 shows the same individual updates we consider above with  $d = 3\Delta x$  (see figure 4-9). The good case is still quite good (bottom right panel), and now the average case also appears quite good (upper right). There are fewer and less severe regions of negative vorticity present, and the analyzed vorticity maxima are about the correct magnitude and about the correct shape. The bad case has remained bad, however. While broadening the vortices has helped alleviate many of the problems seen in the case where  $d = 3\Delta x$ , the initial Lagrangian state vector ensemble (see the left panel of figure 4-5) has such large position discrepancies that some vortices will still be far from truth even when their cores have radius  $6\Delta x$ . As will be made clear in the next chapter, though it is demonstrated here,

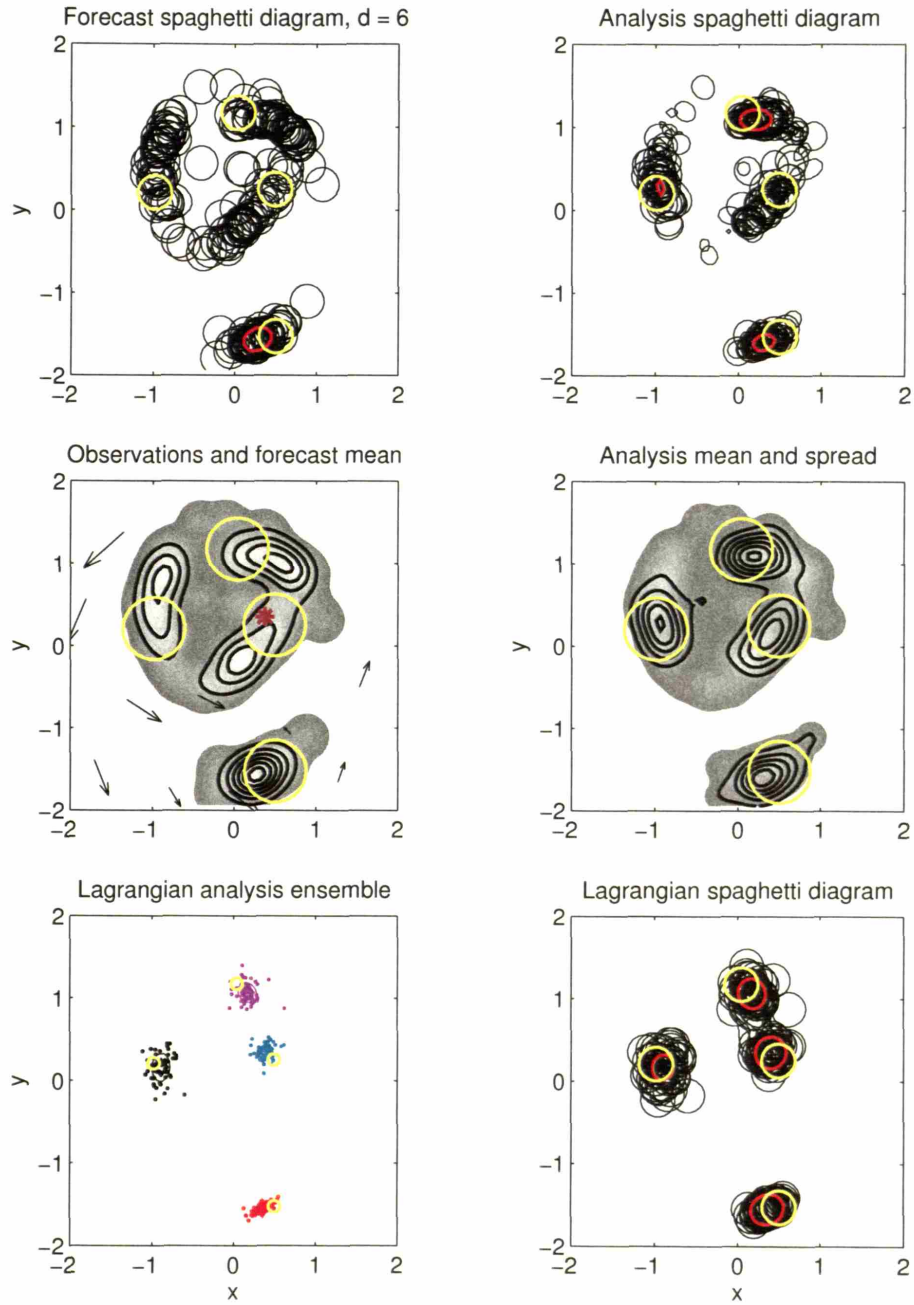


Figure 4-12: These six panels summarize the information shown in figures 4-5 – 4-8, only now for  $d = 6\Delta x$ . The top row shows the prior and analysis ensemble spaghetti diagrams. The contour shown is now one half  $\zeta_{\max}$ . The colors have the same meaning as before. The second row of panels shows the prior and analysis ensemble means and spread. The left panel additionally shows the observations used, ten station velocities and one position. The yellow contours show truth's  $\frac{1}{10}\zeta_{\max}$  contour. The bottom row of panels show the point vortex model updates,  $\mathbf{Z}_L^a$  and  $\mathcal{P}[\mathbf{Z}_L^a; d]$ , respectively.

it is the degree of position dispersion compared to the feature length scale that matters, and the one badly mis-positioned vortex in the middle left panel is very bad, even for this broad vortex scale.

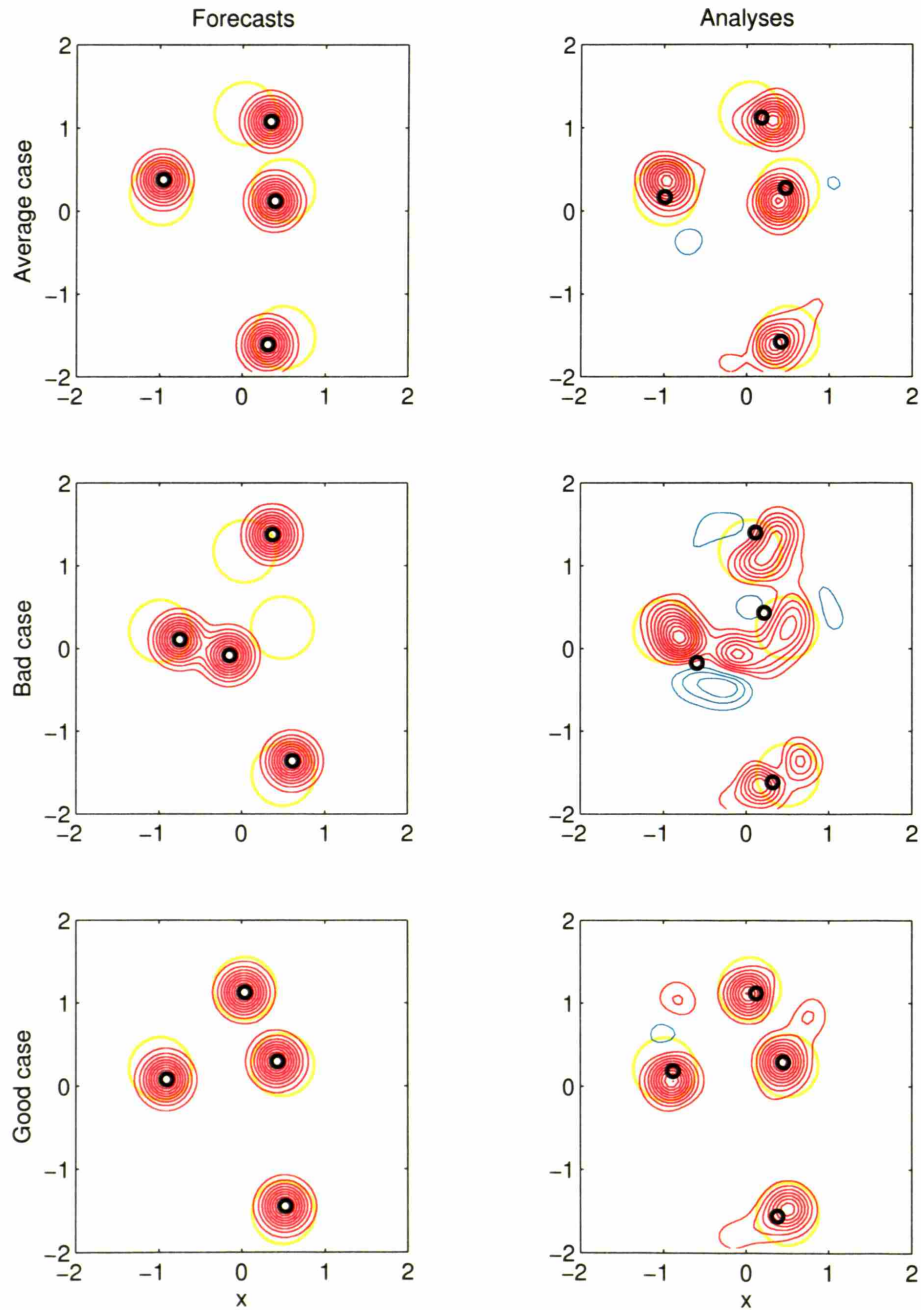


Figure 4-13: This shows the corresponding plots to figure 4-9, only for  $d = 6\Delta x$ . The contour intervals, symbols, and colors are all the same as before.

### 4.2.5 Point vortex summary

In this above, we have seen that ensemble-based Kalman filters can successfully filter Lagrangian model states even though they contain features. If the model state vector is comprised of Lagrangian position information, then a DA scheme can correct position information. However, if a model state vector is comprised of Eulerian field information, then a DA scheme can have trouble correcting mis-positioned features because their associated errors are non-Gaussian when viewed in the Eulerian framework. It does not seem likely that one will be able to transform a general Eulerian problem into a Lagrangian one for estimation purposes, but we have explored the possible benefit of such an approach for a problem where we know it works. The benefit is potentially large, depending on how compact and coherent the features are relative to their dispersion. Unfortunately, as this is not general enough, we must seek an NWP appropriate solution elsewhere. We outline one possible method in chapters 5 and 6 — the development of the approach we present springs from the lessons learned above and from another method of state estimation we explored that also did not prove generalizable. We do not present our work on the matter here, but we suffice it to say we have explored the feasibility of a devising a state estimation scheme modeled after a power curve evolution method called the Level Set Method (e.g., Osher and Sethian 1988; Sethian 1996).

## 4.3 Chapter summary

In this chapter, we have considered position errors and methods to estimate and correct them. NWP models are Eulerian in nature though the phenomena they evolve have apt Lagrangian descriptions. Strong, coherent features like tropical cyclones are common in geophysical fluids, and their treatment inevitably requires specification of their locations. Hence, in predicting the evolution of fluids containing features, one must be mindful of position errors. Though the position errors themselves may be Gaussian, their expression in Eulerian fluid fields like velocity and pressure can easily be rendered non-Gaussian, and this non-Gaussianity violates the basic assumptions of essentially all state-of-the-art DA schemes. This leads to poor state estimates, and in ensemble-based schemes can further lead to physically implausible individual members. To proceed in these scenarios, the state-of-the-art must be improved.

Using a point vortex model, a model defined by the existence of coherent features, configured such that its evolution is nonintegrable, we have seen that state-of-the-art ensemble-based DA schemes can update Gaussian (and weakly non-Gaussian) position errors directly.

This comes as no surprise, though it does suggest a possible solution method for handling features in state estimation is to transform an Eulerian problem into a Lagrangian one. Somewhat obviously, this will never generalize to realistic problems, however, applying this idea to the point vortex model, where we know it does work, we can examine a near best case scenario for the kinds of benefits we can hope to find by considering alternative estimation methods. Also, the point vortex model has given insight to when coherent features should be a problem for state estimation, namely when position errors are comparable the features' length scales. This point will be considered more carefully in the next chapter.

## Chapter 5

# Two-Step Ensemble-Based Data Assimilation

The past three chapters have presented enough background and exploratory material that we are now in a position to approach the state estimation and prediction of spatially extended fluid models that support features. To remind the reader, in chapter 2 we saw how the theoretically correct method for probabilistic filtering is reduced to implementable forms. Nearly all such implementable forms suitable for spatially extended fluid models are forced to assume that a system's error dynamics are linear and retain Gaussian uncertainty distributions. State-of-the-art data assimilation (DA) systems allow extension into regimes of weakly nonlinear error growth via ensemble integrations, but, as explored in chapter 3, they are still limited by these assumptions. Whereas chapter 3 examined generic sources of nonlinear error growth, in chapter 4 we shifted our focus to a specific source, namely that of mispositioned coherent features. We argued that this is a relevant source of error for numerical weather prediction (NWP) scenarios. We further showed that position errors can be, in principle, correctable within a Lagrangian framework by the same state-of-the-art ensemble-based DA schemes that they cause to perform poorly within an Eulerian framework. Stated another way, if a position error is Gaussian in a Lagrangian sense, then it can be seen to satisfy the assumptions of most DA methods *in that Lagrangian sense*, even though position errors do not satisfy those assumptions when viewed within a discretized gridpoint based model. This is encouraging as it suggests we may be able to handle these traditionally difficult scenarios with existing machinery: all that is required is some further consideration of the phenomena at hand. Within this chapter we develop a two-step ensemble-based DA method motivated by the concept of alternative error models (to be defined below). As has been our approach throughout this thesis, we first examine a simpli-

fied example that encapsulates the basic premise; we extend this example to more realistic settings in the next chapter.

Most of the material in this chapter is taken from Lawson and Hansen (2005), and, as such, some of the introductions and motivations may seem redundant.

## 5.1 Introduction

Accurate initial conditions are crucial to the accurate numerical prediction of fluid flows. These initial conditions are the product of state estimation, or data assimilation as it is often called. There is a large literature on methods of state estimation, spanning from a theoretically correct approach to many simplified methods that obtain optimal, or nearly optimal, state estimates when certain assumptions are met. For a variety of practical reasons, we are not able to implement the correct approach in typical geophysical problems, and are thus left to work with the simplifications. Not surprisingly, when the assumptions of these simplified methods are not met, their application can lead to poor estimates, at least locally. Traditional assumptions made are that estimates' errors are (nearly) Gaussian and evolve (nearly) linearly. Another traditional assumption made but not often stated is that estimates' errors are additive in nature. This is the direct consequence of the assumed error model. Most geophysical problems of interest are nonlinear, meaning their errors can readily attain non-Gaussian distributions and possibly cause traditional estimation methods to fail. In some cases where traditional methods fail, assuming error models other than additive ones can satisfy the other assumptions made by these traditional, (nearly) optimal estimation methods, thus allowing their successful application. This chapter examines the notion of alternative error models, and considers the specific example of an error model involving alignment errors.

As discussed in section 2.1.2, filtering is the state estimation technique most relevant to numerical weather prediction (NWP). Filtering seeks the best state estimate given a prior estimate and all the observations available up to the current time, as opposed to smoothing which considers observations further in time than the time at which one wants a state estimate. As no model nor observation is ever perfect, the uncertainties present in prediction problems render filtering an exercise in probability, where the state itself is properly considered a probability density function (PDF).<sup>1</sup> A theoretically correct, fully

---

<sup>1</sup>PDF also sometimes stands for “probability distribution function,” the difference being that density

probabilistic approach to evolving and updating PDFs exists (e.g., Jazwinski 1970), but it is computationally unfeasible to implement, and it requires specification of essentially unobtainable information (e.g., Epstein 1969; Dee 1991). Hence, this theoretically correct approach is often approximated by either moment closure schemes (e.g., stochastic dynamic prediction, Epstein 1969) or Monte Carlo methods like particle filters (e.g., Gordon et al. 1993), both of which must be greatly simplified in order to be implementable by today's computational standards (see 2.3.2 for a discussion of these). Fortunately, much work has been devoted to these simplifications, again with particular attention paid to optimal formulations valid under certain assumptions. No error model assumptions need be made in the full probabilistic approach, but most manageable simplifications must make such assumptions.

For a system with strictly linear dynamics and Gaussian error PDFs, the Kalman filter is the optimal filter (Jazwinski 1970). An optimal estimator is one that minimizes a user-specified error cost function: the two most common optimality criteria obtain the minimum variance estimate and the maximum likelihood estimate, a PDF's conditional mean and mode, respectively (Maybeck 1979). The Kalman filter is a recursive filter that alternates between the evolution and update of a state estimate and its uncertainty, which in a probabilistic framework are properly thought of as the mean vector and covariance matrix of a PDF. The mean and covariance completely determine a Gaussian PDF. The Kalman filter assumes an additive error model, the assumption that deviations between the truth state and estimates thereof can be modeled effectively as a simple additive term, *regardless* of the deviations' physical origin. The Kalman filter assumes the additive term is a realization from a specified, Gaussian random vector distribution, but error models need not assume that errors are Gaussian. *This is an important distinction to draw: an error model is a statement of the degrees of freedom one is allowed to represent and correct state estimates' deviations from truth, not necessarily an assertion of the sources of error in an estimate.* If states are in error because of, say, some gross parametric model flaw (i.e., model inadequacy, or "model error"), then the error model for the update step in a filtering exercise will try to compensate for this flaw in its treatment of the state vector. If the user suspects a parametric flaw, then he or she can of course attempt a parameter estimation problem on the model, but this should be recognized as distinct from an error model's role in state estimation.

Though the Kalman filter is a heavily simplified version of the theoretically correct probabilistic approach, even it remains elusive to implement in the hugely dimensioned problems 

---

functions are presumed to be differentiable.

typical in NWP; indeed, most operational data assimilation schemes to date simplify the Kalman filter further with, for example, assumptions of a static error covariance matrix (Parrish and Derber 1992). Contributing to this difficulty is that many problems scientists wish to filter, particularly in fluid dynamics, are nonlinear. Nonlinearity is a necessary complication to surmount when filtering, for it is nonlinearity that is most often responsible for evolving initially Gaussian PDFs into non-Gaussian PDFs; linear dynamics cannot evolve an initially Gaussian PDF into any other distribution than a differently shaped Gaussian (e.g., circles become ellipses). Out of necessity, various nonlinear extensions to the Kalman filter have been formulated, including ensemble implementations. All of these extensions still rely on the basic framework and machinery of the original linear formulation, most notably its assumed additive error model, however, some assertions of optimality are lost as nonlinear error cost functions may have many local minima. *One can always find a minimum variance estimate regardless of the involved PDFs' forms, but the minimum variance estimate may not always be the estimate of interest to the user, depending on how physically plausible it is.*

When a state estimate's additive errors have become significantly non-Gaussian, one expects poor performance from methods based on the Kalman filter (i.e., most implementable methods); in an ensemble framework, the ensemble mean update can be poor with problems further compounded by unphysical/unrealistic updates of the individual ensemble members comprising the mean. When confronted with these scenarios, instead of trying to implement an estimation method not based on traditional linear/Gaussian methods (i.e., a difficult task), one can try to redefine the problem in terms of linearly correctable errors, that is, errors with Gaussian PDFs. Alternative error models are a natural way to accomplish such a redefinition. We presented a heuristic example of this in section 4.1.1. In figure 4-1, we see that the ensemble of errors (the middle panel) are non-Gaussian and biased, even though we know they are the result of a single unbiased Gaussian random number. One can imagine assuming an alternative error model that allows for an error in the left-right direction (i.e., a simple, constant alignment error); such an alternative error model would allow treatment of the errors as correctable, random realizations from Gaussian PDFs. *This redefinition is the essence of our proposed approach.*

Often an error model will imply an appropriate estimation approach. Extending the heuristic example in section 4.1.1, one can imagine also having additive errors in addition to the alignment errors, thus implying a mixed error model (two different error representations within the same estimate). The notion of mixed error model including alignment errors is treated in some detail below. We demonstrate that this error model naturally

suggests a decomposition of the estimation process into a two-step approach when adhering to linear/Gaussian estimation techniques. Much of this work follows the pioneering study by Hoffman et al. (1995). Guided by a belief that forecast errors should efficiently categorize how forecasts fail, Hoffman et al. argue specifically for an adjustment (originally termed “distortion”) error representation of forecast errors, an error model comprised of alignment (displacement) errors, calibration (amplification) errors, and residual (additive) errors. As error representation is not unique, this work places their approach in the context of alternative error model assumptions. The example error model considered here is similar to Hoffman et al.’s adjustment representation, and as it is physically appealing, its implied two-step approach should be useful for some NWP applications.

This chapter is arranged as follows: section 2 discusses the notion of error models in general and further introduces the specific mixed error model example studied throughout the remainder of the chapter; section 3 examines the implications of alternative error models in state estimation, and considers their application to states from the mixed error model example; section 4 discusses the natural development of a two-step approach and a straightforward approximation applicable when the proper two-step approach is not valid; section 5 presents observation system simulation experiments comparing the performance of the two-step approach and its approximation to a traditional method; and section 6 concludes with a brief summary. *It is important to remember throughout this chapter that the aim is not to determine the best way to do estimation based on alternative error model assumptions, in particular an error model including alignment errors; the aim is to demonstrate the usefulness of alternative error models in state estimation.* Effective implementations of the approach described here suitable for realistic problems with alignment errors will be the subject of the next chapter.

## 5.2 Error models

To understand problems that arise in state estimation procedures, it is important to be clear about the underlying assumptions. Error models can be useful tools in gaining this clarity. For a given dynamical system, consider the truth state,  $\mathbf{x}^t$ , and an estimate of it,  $\mathbf{x}^f$ . For a fluid dynamical system,  $\mathbf{x}^f$  is inevitably an Eulerian estimate represented over a discretized domain. We write the estimate as  $\mathbf{x}^f(s_i)$ , where  $s_i$  denotes the fixed independent coordinate(s) over which a state vector is represented. Unless  $i$  is specified (like  $s_{85}$  in the example in section 4.1.1),  $s_i$  generally refers to all gridpoints, not just one. It is meant to punctuate the state vector’s discrete nature. Here,  $\mathbf{x}^f$  is an  $n$ -vector, meaning that for

$s_i$ ,  $i = 1, n$  for a univariate state vector (if there are, say, 7 different variables defined at each gridpoint, then  $i = 1, n/7$ ). Often truth is a natural phenomenon presumed to be continuous, hence  $\mathbf{x}^t$  should be regarded as the projection of the actual truth into the model space defined by the basis  $s_i$  and how ever many variables are specified at each point. Assume that the estimate deviates from the truth state (i.e., is in error) in some measurable way. The total error in an estimate is often defined as

$$\boldsymbol{\varepsilon}^f \equiv \mathbf{x}^t - \mathbf{x}^f, \quad (5.1)$$

however, there are many ways to imagine representing the origin of the error.<sup>2</sup> That is to say, errors may have more natural representations. How one chooses to represent these deviations from truth is formulated as an error model. Error models are devised with state estimation in mind, and therefore are correctable representations of how estimates are likely to deviate from truth. Error models need not address nor correspond to actual physical sources of error, though it stands to reason that they will be more effective and useful if they do. For example, if a system is driven by multiplicative random noise, an error model with multiplicative errors will probably be most appropriate, but one may well be able to treat the effects of multiplicative noise with a simple additive correction (as traditional Kalman filter-based methods would).

In this section, various error models are introduced to assess the implications of assuming an improper error model in state estimation. Borrowing from the heuristic example in section 4.1.1, we further consider the example of a mixed model including alignment errors; it will be treated throughout the chapter. Though the example is relatively simple in form, it is physically appealing and robust enough to demonstrate the salient points of considering alternative error models.

### 5.2.1 Conceptual error representations

There are many possible error models. Two convenient and often intuitive representations are additive errors and multiplicative errors, also called proportional errors. Respectively, these are:

$$\mathbf{x}^t = \mathbf{x}^f(s_i) + \boldsymbol{\varepsilon}_A(s_i) \quad (5.2)$$

$$\mathbf{x}^t = (\mathbf{1} + \boldsymbol{\varepsilon}_M(s_i)) \circ \mathbf{x}^f(s_i), \quad (5.3)$$

---

<sup>2</sup>The opposite,  $\boldsymbol{\varepsilon}^f \equiv \mathbf{x}^f - \mathbf{x}^t$ , is sometimes used, and is equivalent as long as one is consistent with the definitions of other total errors, for example, observational error is equal to observations less truth instead of vice versa.

where  $\mathbf{1}$  is an  $n$ -vector of 1's and  $\circ$  denotes the Schur product (i.e., element-by-element multiplication; equivalently, one could just as easily construct an  $n$ -by- $n$  identity matrix with  $\mathbf{1} + \boldsymbol{\varepsilon}_M(s_i)$  on the diagonal). Additive error models are probably the most commonly employed error models owing to the definition in equation (5.1), where the error  $\boldsymbol{\varepsilon}_A(s_i)$  is itself an  $n$ -vector added to the estimate to yield a corrected state. Alternatively, multiplicative error models use the  $n$ -vector  $\boldsymbol{\varepsilon}_M(s_i)$  to rescale the estimate at each location, thus achieving its correction. *An additive error model is completely general in that an error vector can always be chosen such that any arbitrary initial estimate is made equal to any arbitrary truth state. A multiplicative error model can be made completely general in the same sense if one takes care to, say, add a constant vector to the initial estimate beforehand to ensure that no element is equal to zero.*

Another intuitive, yet less general, error representation is given by:

$$\mathbf{x}^t = \mathbf{x}^f(s_i + \boldsymbol{\varepsilon}_D(s_i)). \quad (5.4)$$

This error model seeks to alter or displace the basis representation itself until the shape of the initial estimate matches that of the truth state. Note that the error vector is added to  $s_i$  and is itself a function of  $s_i$ . As written, this implies that  $\boldsymbol{\varepsilon}_D$  has a different dimension than  $\boldsymbol{\varepsilon}_A$  or  $\boldsymbol{\varepsilon}_M$  for multivariate state estimates, though one can imagine generalizing this to allow each variable within  $\mathbf{x}^f$  to have its own  $\boldsymbol{\varepsilon}_D$ . *Equation (5.4) implies that an estimation procedure following such an error model would add a correction to the basis rather than the state, thus accomplishing a re-mapping, the simplest scenario being a constant shift of the estimate with respect to the fixed coordinate(s).* If an estimate is displaced in this fashion, information gaps will in general be formed, whether they be at the estimate boundaries or within the estimate in the case that the correction is somehow divergent. Clearly, these are design issues to be addressed by a specific estimation scheme, but they do not prevent consideration of the error model itself. This chapter does not purport how best to accomplish alignment error estimation, it only shows that it can be a useful concept. In the next chapter we discuss qualities we would like alignment error estimation schemes to exhibit, and we offer a plausible approach.

In figure 5-1, we present a schematic to further illustrate the notion that the three error models described by equations (5.2) – (5.4) are all equally applicable for representing the corrective degrees of freedom necessary to improve an estimate to match truth. The top left panel shows a hypothetical state that we will designate “truth,”  $\mathbf{x}^t$ . It has a dimension of  $n = 7$ , and is represented over the basis  $s_i = [1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7]$ . Truth has the value  $\mathbf{x}^t(s_i) = a$  for  $i \neq 5$ , at which point  $\mathbf{x}^t(s_5) = b$ . The top right panel shows a hypothetical

state estimate,  $\mathbf{x}^f$ , of the truth state in the top left panel. The state estimate is very similar to truth except it has  $\mathbf{x}^t(s_3) = b$  and all other elements equal to  $a$ . The three lower panels show the implied error vectors for each of the three error models. The blue curve shows the additive error vector required to correct  $\mathbf{x}^f$  into  $\mathbf{x}^t$ , the green curve shows the multiplicative error vector (as posed in equation (5.3)) required for this correction, and the red curve shows the alignment error vector required for the correction. For the alignment error vector, periodic boundary conditions have been assumed so that all elements are 2 (i.e., just a circular shift), though other forms of  $\boldsymbol{\varepsilon}_A$  are of course available for different boundary conditions. All three error vectors are different from one another, and yet all achieve the desired correction to the truth state.

As is perhaps clear from figure 5-1, an alignment error model is less general than those in equations (5.2) and (5.3) because the degrees of freedom lie with the independent variable. Hence, one can imagine the pathological case where the initial estimate is a constant vector, and thus no amount of displacing its basis representation will lead to any other shape except a constant vector. Yet even with its lack of generality, we maintain that an alignment error model, equation (5.4), is still a useful error model as it is physically appealing. As discussed in section 4.1.1, many classes of problems and phenomena that scientists try to predict exhibit errors of alignment, where some coherent feature is improperly positioned yet properly, or at least reasonably, shaped. One can imagine that while errors at each gridpoint in an estimate can be addressed with an additive error model, coherent Lagrangian structures like tropical cyclones, fronts, and cumulus towers, may be more appropriately addressed by an alignment error model (or more likely, a mixture of both). Section 4.1.2 reviews the historical efforts of atmospheric scientists in dealing with alignment errors, and section 4.1.4 reviews recent efforts in a field that has come to be known as “feature calibration and alignment” (FCA). FCA began with the aforementioned study by Hoffman et al. (1995).

## 5.2.2 Non-uniqueness of error models

As is made clear by figure 5-1, there is no unique way to represent errors. Since additive and multiplicative error models are completely general, and as an alignment error model is sufficient to correct a limited class of estimate errors, combining the degrees of freedom provided by all three error models would seem to give more corrective freedom than necessary, and including the freedom of any other desired error model would affect the process similarly. However, based on an appeal to represent state errors as naturally as possible,

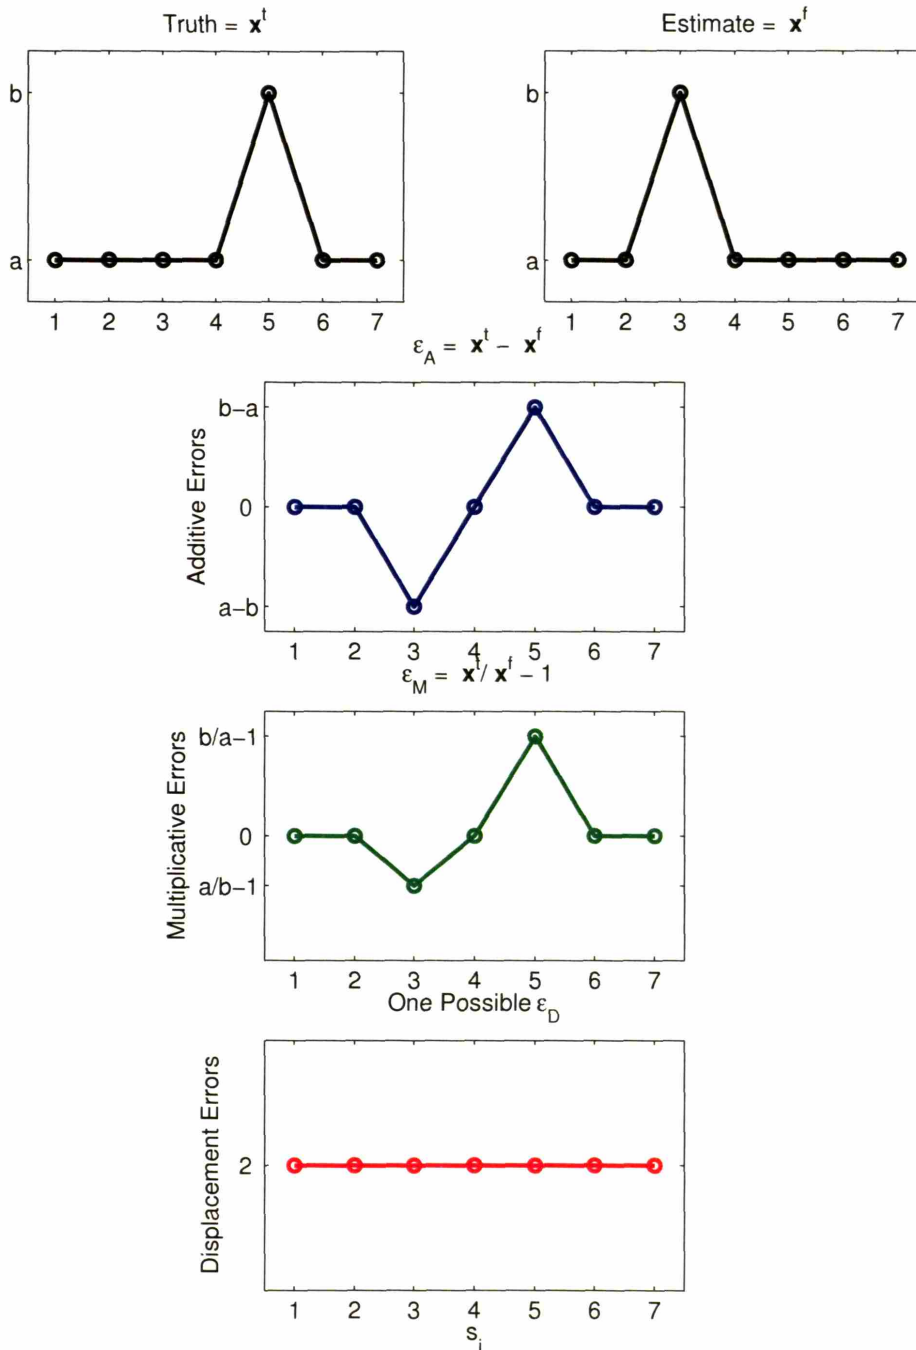


Figure 5-1: The top left panel shows a truth profile, and the top right panel shows an estimate profile. The remaining three panels show different error vectors that can achieve the correction of the estimate profile into the truth profile: the blue curve is for the additive error model, the green curve is for the multiplicative error model, and the red curve is a possible version of an alignment error model.

one may be justified in using such a combined error model. For example,

$$\mathbf{x}^t = (\mathbf{1} + \boldsymbol{\varepsilon}_M(s_i)) \circ \mathbf{x}^f(s_i + \boldsymbol{\varepsilon}_D(s_i)) + \boldsymbol{\varepsilon}_A(s_i) \quad (5.5)$$

would allow estimate improvement through calibration (rescaling,  $\boldsymbol{\varepsilon}_M$ ), alignment (displacement,  $\boldsymbol{\varepsilon}_D$ ), and an additive term ( $\boldsymbol{\varepsilon}_A$ ). This is similar to Hoffman et al.’s (1995) adjustment error representation.

In an ensemble filtering context, there is interest in maintaining an ensemble of states that realistically correspond to truth. As we aim to capitalize on the information given by ensembles while still using implementable linear/Gaussian methods, we conclude that using a combined error model like the one in equation (5.5) is only sensible if the PDF associated with each error vector is nearly Gaussian. *That is to say, using a mixed error model will be useful to state estimation if one can decompose general state errors into the subclasses they fit in the most nearly Gaussian manner, such that the character of any arbitrary  $\boldsymbol{\varepsilon}^f$  could be captured by several, or even many, separately modeled Gaussian random errors.*

### 5.2.3 A mixed alignment and additive error model

This chapter seeks to gain insight into the implications of assuming an improper error model. As error representation is not unique, “improper” here means that errors can be represented more naturally by a different decomposition of Gaussians. Specifically, it asks how well traditional estimation techniques, which assume additive error models, can handle scenarios best described by other error models. To sharpen the focus to a manageable level, we consider an extension to the heuristic example discussed in section 4.1.1, where the total additive error as defined in equation (5.1) is made non-Gaussian by a mixed error model including both an additive error and a simple constant alignment error, that is,  $\boldsymbol{\varepsilon}_D(s_i) = \varepsilon_D$ :

$$\mathbf{x}^t = \mathbf{x}^f(s_i + \varepsilon_D) + \boldsymbol{\varepsilon}_A(s_i), \quad (5.6)$$

where in the heuristic example in section 4.1.1,  $\varepsilon_D$  was assumed to be from an unbiased Gaussian distribution.

This is a particularly nice example to consider as it is physically appealing, with possible applications to correcting and re-positioning some of the common geophysical features discussed in section 4.1.1. *However, it is important to regard this example as one given out of convenience.* Extensions to and complications from “real-world” applications will be discussed in the next chapter, though they do not prevent demonstration of the utility

of alternative error models. In the following subsections, we demonstrate that this error model can produce non-Gaussian total additive errors, and we diagnose when we expect this alternative error model to be of use in state estimation.

### Relating error models

To show the impact of error model assumptions on state estimates, we assume that the mixed error model in equation (5.6) is correct and seek the implications with respect to the Gaussianity assumption within an additive error model. That is, we consider how the errors in the mixed error model in equation (5.6) relate to the total error defined in equation (5.1). Substituting  $\mathbf{x}^t$  from one equation into the other and rearranging the terms yields

$$\boldsymbol{\varepsilon}^f = \mathbf{x}^f(s_i + \varepsilon_D) - \mathbf{x}^f(s_i) + \boldsymbol{\varepsilon}_A(s_i). \quad (5.7)$$

To get an explicit relation between the errors, make a Taylor series expansion of  $\mathbf{x}^f(s_i + \varepsilon_D)$  about  $s_i$ . The leading term in the expansion cancels with the second term on the right hand side of equation (5.7), leaving:

$$\boldsymbol{\varepsilon}^f = \varepsilon_D \frac{d}{ds} \mathbf{x}^f(s_i) + \varepsilon_D^2 \frac{1}{2} \frac{d^2}{ds^2} \mathbf{x}^f(s_i) + \varepsilon_D^3 \frac{1}{3!} \frac{d^3}{ds^3} \mathbf{x}^f(s_i) + \cdots + \boldsymbol{\varepsilon}_A(s_i). \quad (5.8)$$

This equation gives the ideal total error estimate by which a hypothetical filter could additively update the state estimate  $\mathbf{x}^f(s_i)$ , and it turns out that traditional filters cannot achieve this ideal estimate because the PDF of  $\boldsymbol{\varepsilon}^f$  is non-Gaussian in general. To see this, one first needs to specify the expected error statistics of  $\boldsymbol{\varepsilon}_A$  and  $\varepsilon_D$ .

### Expected error statistics

The total error in equation (5.8) depends on the additive error,  $\boldsymbol{\varepsilon}_A$ , and all powers of the alignment error,  $\varepsilon_D$ . For simplicity, we will assume that  $\boldsymbol{\varepsilon}_A$  is a Gaussian random vector from the distribution  $N(\mathbf{0}, \mathbf{P}_A)$ , where  $\mathbf{0}$  is an  $n$ -vector of 0's and  $\mathbf{P}_A$  is the expected covariance matrix of  $\boldsymbol{\varepsilon}_A$ , and  $\varepsilon_D$  is a Gaussian random variable from distribution  $N(0, \sigma_D^2)$ .

The expected values of the alignment error's powers (i.e., its central moments) are:

$$E(\varepsilon_D^N) = \begin{cases} 0 & \text{if } N \text{ is odd} \\ \frac{N!}{(\frac{N}{2})!} \left(\frac{\sigma_D^2}{2}\right)^{\frac{N}{2}} & \text{if } N \text{ is even} \end{cases}$$

where  $E(\cdot)$  denotes expectation (this is the same as the angle brackets that we used back

in chapter 2, e.g., equation (2.1)). This just states the familiar properties of a Gaussian PDF with zero mean — all odd moments are zero and all even moments are determined by the variance, for example,  $E(\varepsilon_D^4) = 3\sigma_D^4$ .

This allows evaluation of the expectations of the various moments of  $\varepsilon^f$  as given in equation (5.8). The expected mean error vector and expected covariance matrix are

$$\boldsymbol{\mu}^f = E(\varepsilon^f) = \frac{\sigma_D^2}{2} \frac{d^2 \mathbf{x}^f}{ds^2} + \frac{\sigma_D^4}{8} \frac{d^4 \mathbf{x}^f}{ds^4} + \dots \quad (5.9)$$

$$\mathbf{P}^f = E(\varepsilon^f \varepsilon^{fT}) = \sigma_D^2 \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^{fT}}{ds} + \frac{\sigma_D^4}{2} \left( \frac{d\mathbf{x}^f}{ds} \frac{d^3 \mathbf{x}^{fT}}{ds^3} + \frac{3}{2} \frac{d^2 \mathbf{x}^f}{ds^2} \frac{d^2 \mathbf{x}^{fT}}{ds^2} + \frac{d^3 \mathbf{x}^f}{ds^3} \frac{d\mathbf{x}^{fT}}{ds} \right) + \dots + \mathbf{P}_A, \quad (5.10)$$

where for simplicity, in equation (5.10) it has been assumed that the two different errors are independent (i.e.,  $E(\varepsilon_D \varepsilon_A) = \mathbf{0}$ ), and superscript T denotes the matrix transpose. If the errors are correlated, then there would be cross-terms as well. Continuing the above expectations on to the third central moment,  $\boldsymbol{\Gamma}^f$ , a third order tensor, one finds that it is in general non-zero, meaning that the total error cannot validly be considered a random variable from a Gaussian PDF. See appendix B for a full derivation of these moments. Also, equation (5.9) shows that the expected mean error is non-zero in general, meaning that the total error is biased. If an error is thought to be biased, it can be taken into account during the estimation process (e.g., Dee and da Silva 1998), but errors are often assumed to have zero mean. Regardless of the inclusion of a bias, traditional filters will still not be able to achieve the ideal additive error estimate given in equation (5.8) because the errors cannot be considered Gaussian — obtaining the ideal error estimate requires information from its higher expected moments.

### Error linearity conditions

If one is able to accurately truncate the Taylor series in equation (5.8) at the term linear in  $\varepsilon_D$ , then the above expected moments for  $\varepsilon^f$  reduce to those of a Gaussian:

$$\begin{aligned} \boldsymbol{\mu}_{\text{lin}}^f &= \mathbf{0} \\ \mathbf{P}_{\text{lin}}^f &= \sigma_D^2 \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^{fT}}{ds} + \mathbf{P}_A \\ \boldsymbol{\Gamma}_{\text{lin}}^f &= \underline{\underline{\mathbf{0}}}, \end{aligned}$$

where  $\underline{\underline{\mathbf{0}}}$  is an  $n$ -by- $n$ -by- $n$  third order tensor of 0's.

To find the general conditions under which we can approximate the Taylor series by its

linear term, first non-dimensionalize the variables so as to compare scales. Assume that  $\mathbf{x}^f \sim X\tilde{\mathbf{x}}$  and  $s \sim L\tilde{s}$  where  $X$  and  $L$  are typical scales with appropriate units multiplying non-dimensional variables of  $O(1)$ , denoted by  $(\tilde{\cdot})$ . This allows comparison of the typical magnitudes of the first and higher terms in the Taylor series. Neglecting the second term in favor of the first requires that their ratio be much less than one. Hence,

$$\frac{\frac{\varepsilon_D^2}{2} \frac{X}{L^2} \frac{d^2 \tilde{\mathbf{x}}}{d\tilde{s}^2}}{\varepsilon_D \frac{X}{L} \frac{d\tilde{\mathbf{x}}}{d\tilde{s}}} = \frac{\varepsilon_D}{2L} \frac{\tilde{\mathbf{x}}_{ss}}{\tilde{\mathbf{x}}_s} \ll 1, \quad (5.11)$$

where subscript  $s$  of  $\tilde{\mathbf{x}}$  denotes differentiation with respect to  $\tilde{s}$ . Since the derivatives are both  $O(1)$  by definition, their coefficient determines whether the inequality holds. The coefficient depends on the expected magnitude of the alignment error and the expected length scale over which the state varies. *Hence, we find that the total error is rendered Gaussian in character by either sufficiently small alignment errors or having relatively broad variations in the state.*

This is a somewhat obvious result, but insofar as the scaling holds, it allows a relatively easy means to identify where non-Gaussianity will be important. For instance, large-scale baroclinic waves are not expected to contribute significantly to this, however, localized features like tropical cyclones might. In fact, in the case of tropical cyclones a typical length scale  $L$  would be the scale of the storm core. This scale itself varies widely from storm to storm, but a representative scale might be  $L \sim 100$  km. An estimate of the alignment error scale can be found from typical hurricane track error scales. DeMaria (1997) shows that typical track errors for 12 hour forecasts are 90 km. These scales make the coefficient in equation (5.11) about one half, which is not much less than one. Hence, tropical cyclones are features whose mis-positioning could violate the assumptions of traditional estimation methods. This comes as no surprise since, as discussed in section 4.1.5, forecasters have found need to implement the messy practice of bogussing precisely to correct these violations (e.g., Leslie and Holland 1995; Surgi et al. 1998).

It should be noted that the condition expressed in equation (5.11) was derived assuming a constant vector of alignment errors,  $\boldsymbol{\varepsilon}_D(s_i) = \boldsymbol{\varepsilon}_D$ , meaning there are no spatial gradients of  $\boldsymbol{\varepsilon}_D(s_i)$ . In the general case where alignment errors are spatially varying, care must be taken in generalizing the scaling argument. However, for the class of alignment errors stemming from mis-positioned coherent features extended over neighboring gridpoints, the spatial gradients of  $\boldsymbol{\varepsilon}_D(s_i)$  are expected to be small so that the features remain coherent (i.e., rapidly varying alignment errors can radically change the shapes of features). Hence, as long as the alignment errors do not vary much more rapidly than the length scale  $L$ , the

same scaling argument should hold locally. That is, if the local value of  $\varepsilon_D$  is comparable to the local spatial gradients associated with any features in the neighborhood, then additive errors can be expected to be non-Gaussian.

Our linearity condition is quite similar in nature to a result from turbulence studies comparing Lagrangian and Eulerian PDFs. There are two general approaches to modeling the statistics of turbulent flows, by considering properties at fixed locations (Eulerian) or by following tagged particles (Lagrangian), and the two approaches produce different PDFs under most circumstances, even when considering the same flow (e.g., Yeung 2002). This makes their comparison difficult, but one would like to be able to compare these PDFs for the purposes of comparing modeling results to theory and observations. Essentially all criteria that have been devised for the conditions under which one can equate the Eulerian and Lagrangian PDFs of a given flow reduce to one sufficient criterion termed the “well-mixed condition” (Thomson 1987; Berloff and McWilliams 2002, Appendix B). An example of a flow satisfying the well-mixed condition is homogeneous, isotropic turbulence, but such a flow is not very general, nor very realistic — as mentioned in section 4.1.1, coherent features often emerge from turbulent flows, particularly rotating ones, and these coherent features can violate the well-mixed condition. In a similar fashion here, violation of our linearity condition implies that Gaussian Lagrangian position errors correspond to non-Gaussian Eulerian amplitude errors. Thus we know of no theories explicitly relating Lagrangian and Eulerian PDFs in our situation of interest, nor can we *a priori* identify situations in which Eulerian non-Gaussianity is associated with Lagrangian Gaussianity beyond using the linearity condition.

We proceed to consider the mixed error model example in detail below. We have highlighted expected problems from applying estimation techniques based on an additive error model to states following the mixed error model. Essentially, all problems arise from having two different Gaussian errors whose cumulative effect leads to total errors that are non-Gaussian. If the alignment errors are small enough, then the total state error is rendered Gaussian in the additive sense. The next section considers the actual estimation of errors and demonstrates the traditional approaches’ deficiencies.

### 5.3 Estimation and error models

Before considering the estimation of states following the mixed error model in equation (5.6), we briefly review the Kalman filter framework and some of its more popular extensions, particularly ensemble methods. Also, in light of the results of the previous section, we

examine the ability of a standard additive error model Kalman filter to reposition a feature. Insight gained will help explain some of the results observed in the previous chapter when an ensemble Kalman filter tried to move a feature (e.g., the point vortices in figure 4-9).

### 5.3.1 Kalman filtering: basic assumptions and machinery

The Kalman filter was introduced and discussed in some detail in section 2.3.3, so we will not repeat that here. But we remind the reader that the filter consists of a forecast step and an update step for what should be considered the mean vector and the covariance matrix of an estimate's PDF. Whether or not the PDFs of concern are actually Gaussian, they are treated as if they were, as the Kalman filter only updates the first two moments. Extensions to higher moments are possible (Miller et al. 1994), but quickly become exorbitantly expensive for large-dimensional systems.

The forecast step of the Kalman filter evolves a state vector estimate and its uncertainty forward in time under some assumed system dynamics (i.e., a model). If the model is linear, then the state vector estimate and its covariance matrix can be evolved with the same model. If the model is nonlinear, then the uncertainty must be evolved approximately. One approach, the extended Kalman filter (EKF, Jazwinski 1970; Ghil and Malanotte-Rizzoli 1991), uses the dynamics' tangent linear model evaluated about the nonlinearly evolved state vector estimate forecast for this purpose. This approach is subject to the validity of its tangent linearity assumption and is generally computationally expensive. Ensemble methods are another way to approximately evolve a state's uncertainty. We will apply both of these filters to a specific example below.

To remind the reader, the state vector update equation is traditionally written as

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{K}(\mathbf{y}^o - \mathbf{H}\mathbf{x}^f), \quad (5.12)$$

where

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^T (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1}, \quad (5.13)$$

and  $\mathbf{x}^f$  is a prior state estimate (an  $n$ -vector),  $\mathbf{P}^f$  is the prior estimate's error covariance (an  $n$ -by- $n$  matrix),  $\mathbf{y}^o$  is an  $m$ -vector holding the available observations,  $\mathbf{R}$  is the observations' error covariance (an  $m$ -by- $m$  matrix), and  $\mathbf{x}^a$  is the updated state vector estimate, called the analysis (the analysis error covariance matrix can also be found). As seen in equation (5.12), the analysis is a specially weighted linear combination of  $\mathbf{x}^f$  and  $\mathbf{y}^o$ , with the weighting determined by the Kalman gain matrix,  $\mathbf{K}$  (an  $n$ -by- $m$  matrix), which compares the two information sources' relative uncertainties. Since  $m$  is often (much) less than  $n$ ,

an observation operator,  $\mathbf{H}$ , is presumed to exist that transforms information in  $\mathbf{x}^f$  to a form that is meaningfully comparable to  $\mathbf{y}^o$ . The differences between the two information sources,  $(\mathbf{y}^o - \mathbf{H}\mathbf{x}^f)$ , termed the innovation vector, determine how to improve the prior estimate to arrive at the analysis. Together,  $\mathbf{K}$  and the innovation vector comprise the analysis increment, a fully  $n$ -dimensioned, additive correction term in accordance with the assumed error model in equation (5.2). Note that the observation operator and its matrix transpose are necessary to meaningfully compare the uncertainties in the two different covariance matrices. If the observation operator is nonlinear, denoted by  $\mathcal{H}[\cdot]$ , then the full operator is used within the innovation vector, but its Jacobian matrix evaluated about the current state,  $\mathbf{H}$ , is used to form  $\mathbf{K}$ . We discuss the subtleties of using nonlinear observation operators in the next section.

### 5.3.2 Ensemble-based filtering

Though Monte Carlo methods can be used to approximate the correct, fully nonlinear probabilistic approach, the ensemble methods considered here are those that facilitate the implementation of the Kalman filter machinery in systems with nonlinear dynamics. As such, it is useful to state the goals of this breed of ensemble filtering (Anderson and Anderson 1999; Anderson 2001). The first goal is to produce an ensemble mean with small errors — as these ensemble methods are rooted in assumptions of nearly linear error dynamics with nearly Gaussian PDFs, the ensemble mean should often approximate the PDF’s mode, the most likely realization from the PDF and hence its “best estimate.” As error representation is not unique, “small errors” typically means the root mean square (rms) of the total (additive) error is in some sense small. The second goal is that all individual ensemble members should be statistically indistinguishable from truth — effective Monte Carlo approximation should use an ensemble comprised of random draws from the same PDF that produced truth.

Evensen (1994) was the first to use an ensemble of forecasts to estimate the PDFs necessary in Kalman filtering. In Evensen’s so-called ensemble Kalman filter (EnKF), the key assumption is equating the covariance of an ensemble of  $N$  forecasts, denoted by  $\mathbf{z}_j^f$ , with  $j = 1, N$ , each evolved with the full nonlinear model, with the linearly evolved forecast error covariance matrix necessary for the EKF. Concomitantly, one must identify the ensemble mean as one’s best state estimate. As such, traditional ensemble-based methods, including the EnKF (e.g., Houtekamer and Mitchell 1998; Burgers et al. 1998) and its so-called “deterministic” cousins (e.g., Bishop et al. 2001; Anderson 2001; Whitaker and Hamill 2002), are designed to update the ensemble mean, in accordance with the first goal of ensemble filter-

ing. Consequently, the mean is updated at the expense of the individual ensemble members — the ensemble members are updated to enforce their resulting mean being the desired approximate minimum variance estimate. This works well when the individual ensemble members do not substantially differ from the mean, however, as ensemble spread becomes large, individual members can have completely unphysical updates, thereby violating the second goal of ensemble filtering.

*One important advantage of ensemble methods is that they require no explicit linearizations in either the system dynamics or the observation operator.* Making  $N$  separate non-linear ensemble integrations allows the states to disperse in whatever error directions are important, thus allowing direct approximation of the expected error covariance matrix of the total error, regardless of how non-Gaussian its actual PDF may be. Similarly, as first suggested by Evensen (1994), and as explained in section 2.3.3, since one is using the ensemble to estimate  $\mathbf{P}^f$ , one can just as easily use the ensemble to estimate  $\mathbf{P}^f \mathbf{H}^T$  and  $\mathbf{H} \mathbf{P}^f \mathbf{H}^T$  directly by finding the population covariance of the ensemble,  $\mathbf{z}_j^f$ , with the “observed ensemble”,  $\mathcal{H}[\mathbf{z}_j^f]$ , or the covariance between the observed ensemble and itself.

Direct use of full observation operators enables ensemble methods to include special observations in cases where the typical extended/linearized methods cannot. Most relevantly, in situations where strong coherent features like tropical cyclones are present, special position observations are often taken. As explained in section 4.2.4, an observation operator that elicits the position of a feature from a gridpoint-based state vector will effectually evaluate the function  $\arg \max[\mathbf{x}(s_i)]$  (sometimes also called “*max loc*”), the function that elicits the position  $s_{max}$  from the state  $\mathbf{x}(s_i)$ . The Jacobian matrix of this function, that is, the linear sensitivity of  $s_{max}$  to the values of the state’s elements, is either zero or  $\pm$  infinity. This is essentially a statement that the state can determine a feature’s position, but a feature’s position cannot linearly affect the state. This is also true when *arg max* is applied to the state’s spatial gradient. Other approaches to finding  $s_{max}$  are possible, such as local interpolation using the output from *arg max* as a seed or fitting a polynomial of order  $n - 1$  to the gridded data, differentiating it, and solving for the zero(s), but none have yet proven feasible to linearize, meaning that innovations in position information cannot affect the state vector estimate. Because ensemble methods avoid evaluating the Jacobian matrix of  $\mathcal{H}[\mathbf{x}(s_i)]$ , they allow position innovations to correct state estimates; however, one must still be aware of the inconsistencies explained in section 2.3.3 (below equation (2.43)) that can arise in the linear combination of the update step.

Applying ensemble-based filters to estimation problems with Gaussian errors in regimes of near-linear error growth works very well and satisfies both stated goals of ensemble

filtering, that is, that the ensemble mean has small rms errors and each ensemble member be statistically indistinguishable from truth. In chapter 3 we found that as the error growth becomes appreciably nonlinear, thus allowing errors to become non-Gaussian, these methods work less effectively. Hence, they are not expected to work effectively when applying them to states following the mixed error model outside of its linear range as defined above in section 5.2.3. This is mainly because dispersed feature positions will produce a mean that is smoothed and thus unlike any of the ensemble members that comprise it. To see these deficiencies in execution, we will return to the mixed error model introduced in the previous section where synthetic update examples of an ensemble are examined. However, before examining these update examples, we first examine the Kalman filter’s general ability to reposition coherent features.

### 5.3.3 The Kalman filter as a repositioning device

In this subsection we examine the ability of the Kalman filter machinery to move a feature within a domain. This builds from the development begun in section 5.2. To remind the readers, we are considering the state estimate  $\mathbf{x}^f$ , an  $n$ -vector, of a truth state,  $\mathbf{x}^t$ . The physical representation of  $\mathbf{x}^f$  is quite general, it can be univariate or multivariate, and it can occupy one, two, or three spatial dimensions. Forecast models require knowledge of the state vector’s physical representation, but DA systems like the Kalman filter require only the state vector (written as a vector, regardless of whether the domain is two or three-dimensional) and its covariance matrix. All physical knowledge of the system for the purposes of state estimation is contained within the covariance matrix, and though this is statistical, it is potentially a large amount of information. Indeed, as mentioned in section 2.3.3, Dee (1991) argues that it is too much information to specify meaningfully for most estimation problems. However, we have indicated above in section 5.2 that the physical knowledge held by the covariance matrix is not sufficient for describing and constraining mis-positioned coherent features — their proper correction requires physical knowledge held by higher order statistical moments. What are the consequences of DA schemes like the Kalman filter discarding these important higher-order moments?

To treat alignment errors, we explicitly introduced the state vector’s basis representation,  $s_i$ , such that  $\mathbf{x}^f = \mathbf{x}^f(s_i)$ . The basis is written with a subscript to emphasize its discrete nature. Explicitly writing the basis as a separately indexed quantity is meant to suggest treating the state vector values as the dependent variable (the ordinate) and the basis values as the independent variable (the abscissa). Even if this is not a strictly valid interpretation, in some senses this is what numerical models do: gridpoint locations are often fixed in time,

but even if they are not (e.g., adaptable meshes), they are among the only quantities of which models are ever fully confident. However, one way to address position error is to treat the gridpoint locations as if they were uncertain, thereby allowing one to move the location of a state element with a given value. The philosophy and mechanics of this realigning will be addressed in the next chapter.

In section 5.2, we found that if a state estimate has a simple, Gaussian, constant alignment error (i.e., a left or right shift), then its corresponding additive error representation is generally non-Gaussian. If we neglect the additive error component in equation (5.8), then we are left with the following expression for the additive error due to an alignment error:

$$\boldsymbol{\varepsilon}^f(s_i) = \varepsilon_D \frac{d}{ds} \mathbf{x}^f(s_i) + \varepsilon_D^2 \frac{1}{2} \frac{d^2}{ds^2} \mathbf{x}^f(s_i) + \varepsilon_D^3 \frac{1}{3!} \frac{d^3}{ds^3} \mathbf{x}^f(s_i) + \dots \quad (5.14)$$

Errors are random realizations from PDFs, whereas analysis increments are *estimates* of errors. To distinguish these related quantities, we generally denote errors by  $\varepsilon$  and increments by  $\Delta$ . This means that the Kalman filter update equation (see equation (5.12)) can be written as  $\mathbf{x}^a = \mathbf{x}^f + \Delta \mathbf{x}$ , where  $\Delta \mathbf{x} = \mathbf{K}(\mathbf{y}^o - \mathcal{H}[\mathbf{x}^f])$ . Comparing this to the additive error model written in equation (5.2), we see that  $\Delta \mathbf{x}$  is the filter's best estimate of  $\boldsymbol{\varepsilon}^f$ . Therefore, to correct the alignment error,  $\varepsilon_D$ , a hypothetically ideal DA scheme that assumes an additive error model would find the analysis increment:

$$\Delta \mathbf{x}_{\text{ideal}} = \Delta s \frac{d\mathbf{x}^f}{ds} + \Delta s^2 \frac{1}{2} \frac{d^2 \mathbf{x}^f}{ds^2} + \Delta s^3 \frac{1}{3!} \frac{d^3 \mathbf{x}^f}{ds^3} + \dots, \quad (5.15)$$

where  $\Delta s$  is the best estimate of the position error. For a sufficiently accurate  $\Delta s$ , this increment would be able to reposition the state vector. The question is then how close a Kalman filter could come to this ideal analysis increment. We know from section 5.2 that the Kalman filter's increment will differ, but we seek to understand how it will differ.

To gain insight to the way the Kalman filter would correct an alignment error, we examine the only clean example possible, the case where we know the Kalman filter will work. As explained above at the end of section 5.2, the Kalman filter will work when the Taylor series in equation (5.14) can be truncated at its linear term, meaning  $\varepsilon_D$  is small in the sense explained in the same section. Hence, we suppose:

$$\boldsymbol{\varepsilon}^f(s_i) = \varepsilon_D \frac{d}{ds} \mathbf{x}^f(s_i), \quad (5.16)$$

which implies that:

$$\Delta \mathbf{x}_{\text{ideal}} = \Delta s \frac{d}{ds} \mathbf{x}^f(s_i). \quad (5.17)$$

If an observation is taken that is much more accurate than the prior estimate, then the Kalman filter will produce an analysis that agrees almost exactly with the observation. Hence, if we assume a single, perfect observation is available at the location  $s_o$ , then we can examine how the Kalman filter re-positions the state estimate to agree with the observation, and hence deduce the Kalman filter's estimate of  $\Delta s$ . We write the observation as  $\mathbf{y}^o = y_o$  and the prior estimate as  $\mathcal{H}[\mathbf{x}^f(s_i)] = \mathbf{x}^f(s_o) = x_o$ . The Kalman gain matrix is defined as  $\mathbf{K} = \mathbf{P}\mathbf{H}^T(\mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R})^{-1}$  (though in this one-dimensional case  $\mathbf{K}$  is a column vector). The case of a single observation means that  $\mathbf{H}\mathbf{P}\mathbf{H}^T$  and  $\mathbf{R}$  are scalars, and the case where the observation is essentially perfect means that  $\mathbf{R} \rightarrow 0$ . Hence, the gain matrix simplifies to:

$$\mathbf{K} = \frac{1}{\mathbf{H}\mathbf{P}\mathbf{H}^T} \mathbf{P}\mathbf{H}^T.$$

We can evaluate these terms by directly taking the expectations from equation (5.16):

$$\begin{aligned} \mathbf{P}\mathbf{H}^T &= \mathbb{E}(\boldsymbol{\varepsilon}^f(s_i)\boldsymbol{\varepsilon}^{fT}(s_o)) & \mathbf{H}\mathbf{P}\mathbf{H}^T &= \mathbb{E}(\boldsymbol{\varepsilon}^f(s_o)\boldsymbol{\varepsilon}^{fT}(s_o)) \\ &= \mathbb{E}(\varepsilon_D^2) \left. \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^f}{ds} \right|_{s_o} & &= \mathbb{E}(\varepsilon_D^2) \left( \left. \frac{d\mathbf{x}^f}{ds} \right|_{s_o} \right)^2 \\ &= \sigma_D^2 \left. \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^f}{ds} \right|_{s_o} & &= \sigma_D^2 \left( \left. \frac{d\mathbf{x}^f}{ds} \right|_{s_o} \right)^2, \end{aligned}$$

where  $\sigma_D$  is the expected size of  $\varepsilon_D$  (i.e., the error standard deviation). This gives a gain matrix:

$$\mathbf{K} = \frac{d\mathbf{x}^f}{ds} \left( \left. \frac{d\mathbf{x}^f}{ds} \right|_{s_o} \right)^{-1}. \quad (5.18)$$

Note that in this linear example,  $\sigma_D^2$  drops out of the expression.

Comparing equation (5.17) to the expression for the analysis increment,  $\Delta \mathbf{x} = \mathbf{K}(y_o - x_o)$ , we immediately identify

$$\Delta s = \left( \left. \frac{d\mathbf{x}^f}{ds} \right|_{s_o} \right)^{-1} (y_o - x_o). \quad (5.19)$$

Hence, the ideal analysis increment,  $\Delta \mathbf{x}_{\text{ideal}}$ , is just the first spatial derivative of the estimate scaled by a factor proportional to the estimate and truth's difference (i.e., the innovation). From a dimensional (here "dimensional" referring to units) standpoint, the innovation is in units of  $\mathbf{x}$ , whereas  $\Delta s$  is in units of  $s_i$ , so the factor of the inverse derivative is required to make this dimensional conversion. Pre-factor aside,  $\Delta s$  is directly proportional to the size of the innovation. For coherent features with sharp spatial gradients, a small  $\varepsilon_D$  can cause large values of  $\boldsymbol{\varepsilon}^f$ , at least locally. This in turn can give a large innovation, on the order of

the feature's amplitude itself.

Without going through the detailed derivation, one can see that the expression for  $\mathbf{K}$  becomes less clean as more terms are retained in the Taylor series in equation (5.14). For example, the expression for the gain matrix when retaining two terms is:

$$\mathbf{K} = \left( \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^f}{ds} \Big|_{s_o} + \frac{3}{4} \sigma_D^2 \frac{d^2\mathbf{x}^f}{ds^2} \frac{d^2\mathbf{x}^f}{ds^2} \Big|_{s_o} \right) \left( \left( \frac{d\mathbf{x}^f}{ds} \Big|_{s_o} \right)^2 + \frac{3}{4} \sigma_D^2 \left( \frac{d^2\mathbf{x}^f}{ds^2} \Big|_{s_o} \right)^2 \right)^{-1},$$

and now cleanly identifying a  $\Delta s$  from  $\mathbf{K}(y_o - x_o)$  is not possible. Contributing to this is the fact that retaining two terms in the Taylor series in equation (5.14) gives an error bias and a non-negligible third central moment, as described in section 5.2 and appendix B, respectively. Also, whereas the expression for the ideal increment in equation (5.15) uses successive powers of  $\Delta s$  itself, the Kalman filter increment relies heavily on  $\sigma_D$ , the expected error size, rather than the best estimate of the error. *Therefore, coherent features pose at least two serious problems to the Kalman filter: the first is that the Kalman filter cannot obtain the correct form of the gain matrix due to neglecting higher order moments (and the error bias), and the second is that mis-positioned coherent features can give quite large innovations, and these large innovations will amplify an already incorrect gain matrix, making existing problems worse.*

So what has this taught us? One of the most important lessons to take away from this is that the Kalman filter is relatively restricted when it comes to trying to realign estimates. Most of its effective re-alignment comes from adding scaled versions of a profile's first spatial derivative. If it must correct an  $\varepsilon_D$  comparable to a feature's length scale (as described above in section 5.2) then adding the first derivative scaled by a large number can lead to deformities in the analysis profile. This phenomenon helps explain the emergence of the "pools" of negative vorticity in the point vortex examples shown in section 4.2.4, figure 4-9. A one-dimensional schematic example of this is shown in figure 5-2. The top panel shows a hypothetical base profile with a feature in it and its first spatial derivative. As in the point vortex example in section 4.2.1, we choose a Gaussian "hump" with form  $\mathbf{x} = A \exp(-Bs^2)$  as the base profile; therefore,  $\frac{d\mathbf{x}}{ds} = -2Bs\mathbf{x}$ . To confirm that most of the re-alignment comes from scaled versions of  $\frac{d\mathbf{x}}{ds}$ , the middle panel shows this quantity along with 16 possible gain matrix forms. Each light blue line is a column from the covariance matrix one gets assuming that  $\varepsilon_D$  is very small (i.e., where we expect the Kalman filter to work). The covariance matrix has been found from a very large ensemble. The columns chosen for display correspond to the  $\mathbf{PH}^T$  that would be generated for each of 16 different possible station observation locations, all locations being slightly to the left of the base

profile's peak. Each line here would simply be scaled by the amount  $(\frac{dx}{ds}|_{s_o})^{-1}(y_o - x_o)$  to give an analysis increment, and it is evident that all are essentially of the same form as  $\frac{dx}{ds}$ . The bottom panel shows a similar plot to the middle panel, only the covariance matrix was generated assuming that  $\varepsilon_D$  was moderately large, about half the Gaussian's length scale (which we take to be its half-width). This is a case where the Kalman filter is not expected to reposition the feature correctly. Analysis increments made from these  $\mathbf{PH}^T$  can cause generally undesirable effects like forming regions of negative  $\mathbf{x}$ , strongly increasing the Gaussian hump's amplitude, and possibly even forming multiple maxima. The forms of the  $\mathbf{PH}^T$  vectors show that they differ from  $\frac{dx}{ds}$  (given by the red dashed line), but not hugely. Though this is for the very special case with constant alignment errors and a single perfect observation, experience has shown it to be a useful model for explaining many of the problems found in more complicated, general cases.

We now return to the mixed error model introduced in the previous section where synthetic update examples of an ensemble are examined.

### 5.3.4 Examples of filter updates

When considering examples, there are two separate roles assumed, and it is important to differentiate between them. First, the creator designs the problem; this role is in a sense omniscient. Second, the discoverer tackles the problem set up by the creator; the discoverer enters a problem armed only with limited information and assumptions typical to traditional state estimation. Obviously, if the discoverer knew what the creator knew, then these estimation examples would be trivial. As with all examples in this thesis, a perfect model scenario is assumed throughout these experiments to clarify the underlying update mechanics. This decision is also guided by a belief that one must understand the perfect model scenario before being able to understand how known or unknown sources of model inadequacy will complicate matters.

#### The Korteweg-de Vries equation

The Korteweg-de Vries (KdV) equation is a well studied nonlinear partial differential equation. The model dynamics with all coefficients set to 1 is:

$$\mathbf{u}_t + \mathbf{u}\mathbf{u}_s + \mathbf{u}_{sss} = \mathbf{0}, \tag{5.20}$$

where  $\mathbf{u} = \mathbf{u}(s, t)$  and subscripts denote differentiation with respect to that variable. One remarkable aspect of the KdV equation is that it admits specific, coherent analytical solu-

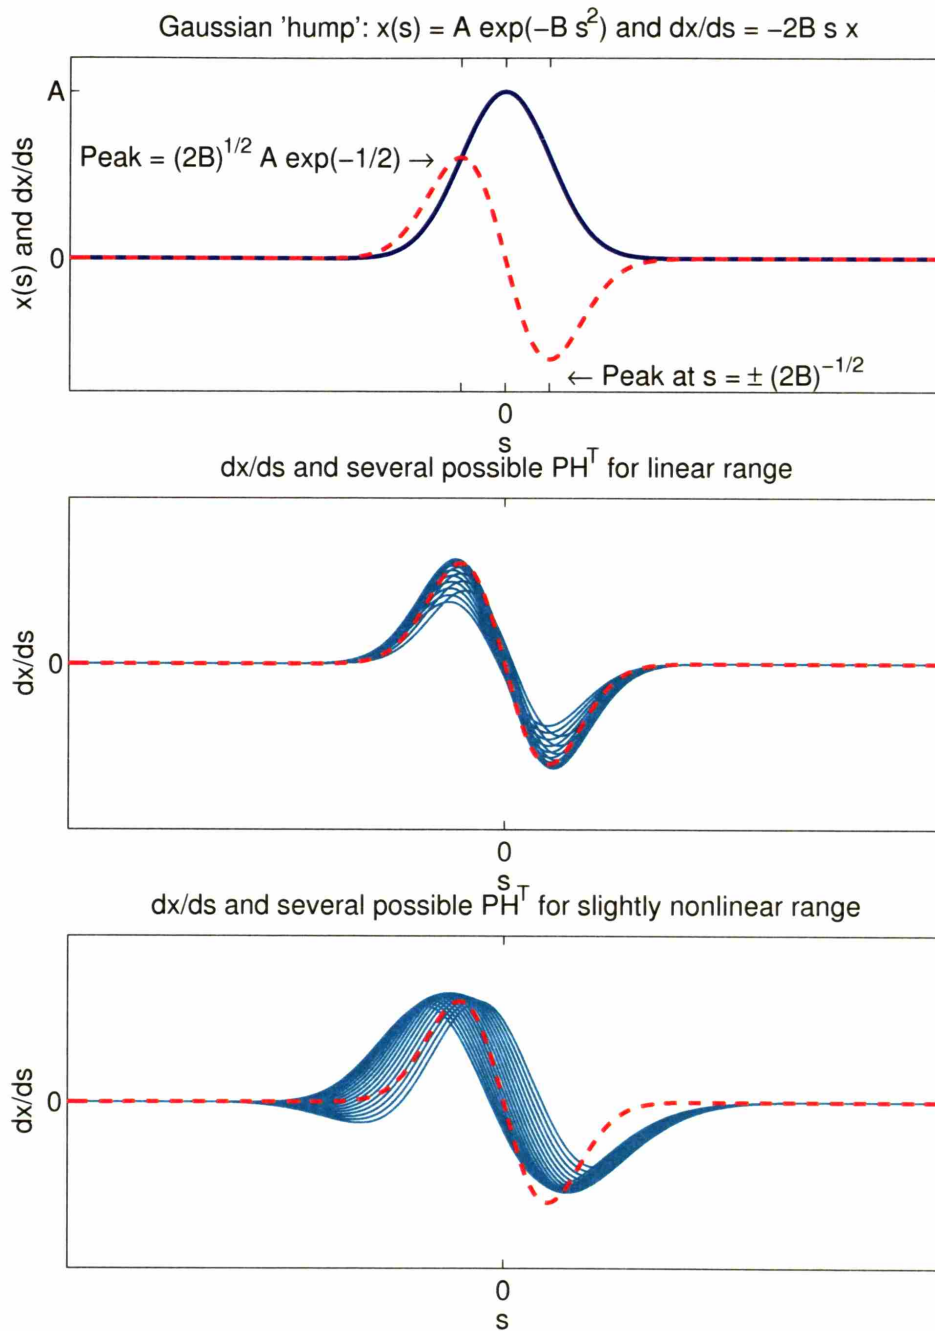


Figure 5-2: The top panel shows a Gaussian “hump” profile,  $\mathbf{x}(s) = A \exp(-Bs^2)$ , and its first derivative. The middle and bottom panels compare the first derivative profile to the profiles of selected columns in the covariance matrix  $\mathbf{P}$ . Each column of  $\mathbf{P}$  corresponds to the  $\mathbf{PH}^T$  vector that would arise from a single given station observation. The middle panel is taken from a  $\mathbf{P}$  generated from small  $\varepsilon_D$ , and the bottom panel assumes a moderately sized  $\varepsilon_D$ .

tions called solitons:

$$\mathbf{u}_{\text{sol}}(s, t) = 3A \operatorname{sech}^2 \left( \frac{A^{1/2}}{2} (s - At) \right), \quad (5.21)$$

where  $A$  is some constant. Solitons are elastic wave-forms that propagate with speeds proportional to their amplitudes. The KdV equation admits many other solutions with no known analytical expressions. It is important to note that the equation is stable to small perturbations (Drazin and Johnson 1989). Here the KdV equation is not treated as a realistic model of fluid motion, but rather as a dynamical system which supports “features”.<sup>3</sup>

This provides a useful testbed for filtering and prediction. The creator knows that truth is a soliton expressible in the form of equation (5.21), but the discoverer only knows that the model dynamics follow equation (5.20). Periodic boundary conditions are assumed. Note that as the propagation speed and amplitude are related through the constant  $A$ , errors in position and additive errors are not likely to be strictly separable. The examples in this section consider only the update step of filtering. The examples shown in section 5 include the forecast step.

The KdV equation has been used before in estimation and filtering contexts. Muccino and Bennett (2002) assumed that the KdV equation was the governing dynamics for internal solitary waves observed in the ocean, and used realistic observations to perform a generalized inversion of the model, allowing for uncertainty in the model dynamics as well. Recently, van Leeuwen (2003) has used the KdV equation in a filtering exercise to test his particle filter-based data assimilation system, the sequential importance resampling filter (SIRF, also see section 2.3.2).

## A look at total errors

The statistics of the total errors are key to estimation success. The results of section 5.2 indicate that the total errors can easily be non-Gaussian and biased. This has already been presented in the heuristic example in section 4.1.1. The base profile in the top panel of figure 4-1 is in fact the soliton profile,  $\mathbf{u}_{\text{sol}}$ , in equation (5.21), where the creator has chosen  $A = 1$  (consider  $t = 0$ ), a domain length of 40, and  $n = 200$ , giving a gridpoint spacing of 0.2. Note that the “feature” has a half-width of 3.5.

We form an ensemble in accordance with the mixed error model, equation (5.6), by the realization of two random, uncorrelated variables,  $\varepsilon_D$  and  $\varepsilon_A$ . For the purposes of clear display, the creator chooses to have the additive errors perfectly and uniformly covary

---

<sup>3</sup>When applied to an actual hydrodynamical fluid system, there are coefficients for the terms that depend on gravity and some medium-dependent constants.

over the domain length, meaning that the covariance matrix of  $\boldsymbol{\varepsilon}_A$  is singular:  $\mathbf{P}_A = \sigma_A^2 \mathbf{1}$ , where  $\mathbf{1}$  is an  $n$ -by- $n$  matrix of 1's. This amounts to simply shifting the entire profile up or down — this does not affect the generality of using a full rank  $\mathbf{P}_A$  (for a discussion of singular covariance matrices, see the last subsection in section 2.3.3). Creator chooses the error statistics to be as follows:  $\alpha_j$  is a realization of  $\varepsilon_D$  from the PDF  $N(0, \sigma_D^2)$ , where  $\sigma_D = 1.75$  to ensure position dispersion comparable to the feature's length scale;  $\beta_j$  is a realization that determines  $\boldsymbol{\varepsilon}_A$  and is from the PDF  $N(0, \sigma_A^2)$ , where  $\sigma_A = 0.15$ , 5% of the total amplitude. Setting the base profile,  $\mathbf{x}^b$ , equal to  $\mathbf{u}_{\text{sol}}$ , ensemble members are generated by:

$$\mathbf{z}_j(s_i) = \mathbf{x}^b(s_i + \alpha_j) + \beta_j \mathbf{1}, \quad \text{for } j = 1, N, \quad (5.22)$$

where  $j$  is the ensemble member index and  $N$  is the total number of ensemble members.

Figure 5-3 shows an ensemble realized from equation (5.22) and the error statistics given in the paragraph preceding it. This figure shows the same information as shown in figure 4-1, only now  $\beta_j \neq 0$ . The light blue lines in the top panel show a 50 member ensemble of shifted soliton states. The blue dotted line is their ensemble mean (one dot per gridpoint), and the black dashed line is the base soliton profile from which the ensemble was formed. The second panel shows the ensemble of total errors formed from subtracting each ensemble member from the base profile, and the blue dotted line is again their mean. The vertical dotted line shows the location at which a histogram was compiled for the values of the total errors. The bottom panel shows a normalized histogram compiled from a similarly formed 1000 member ensemble.

For the purpose of calculating the expected statistics of the total errors, we use the expected value of truth, that is,  $\alpha_j = \beta_j = 0$ . Note that ensemble methods typically assume the ensemble mean is the expected value of truth, but due to the soliton position dispersion, this is not the case here, hence, there is a bias in the total error statistics (see the non-zero mean in the middle panel of figure 5-3). Using  $\mathbf{x}^t = \mathbf{x}^b (= \mathbf{u}_{\text{sol}})$ , one can evaluate the total error of each ensemble member:

$$\boldsymbol{\varepsilon}_{\mathbf{z}_j} = \mathbf{x}^t - \mathbf{z}_j, \quad \text{for } j = 1, N. \quad (5.23)$$

In the limit of very large  $N$ , the PDF approximated by this ensemble of errors should approach the expected PDF of  $\boldsymbol{\varepsilon}^f$ , the total error as defined in equation (5.1) and made explicit for this mixed error model in equation (5.8).

These errors are significantly non-Gaussian. This can be confirmed by using multivariate estimates of population skewness and kurtosis (Mardia 1970), though these can be difficult

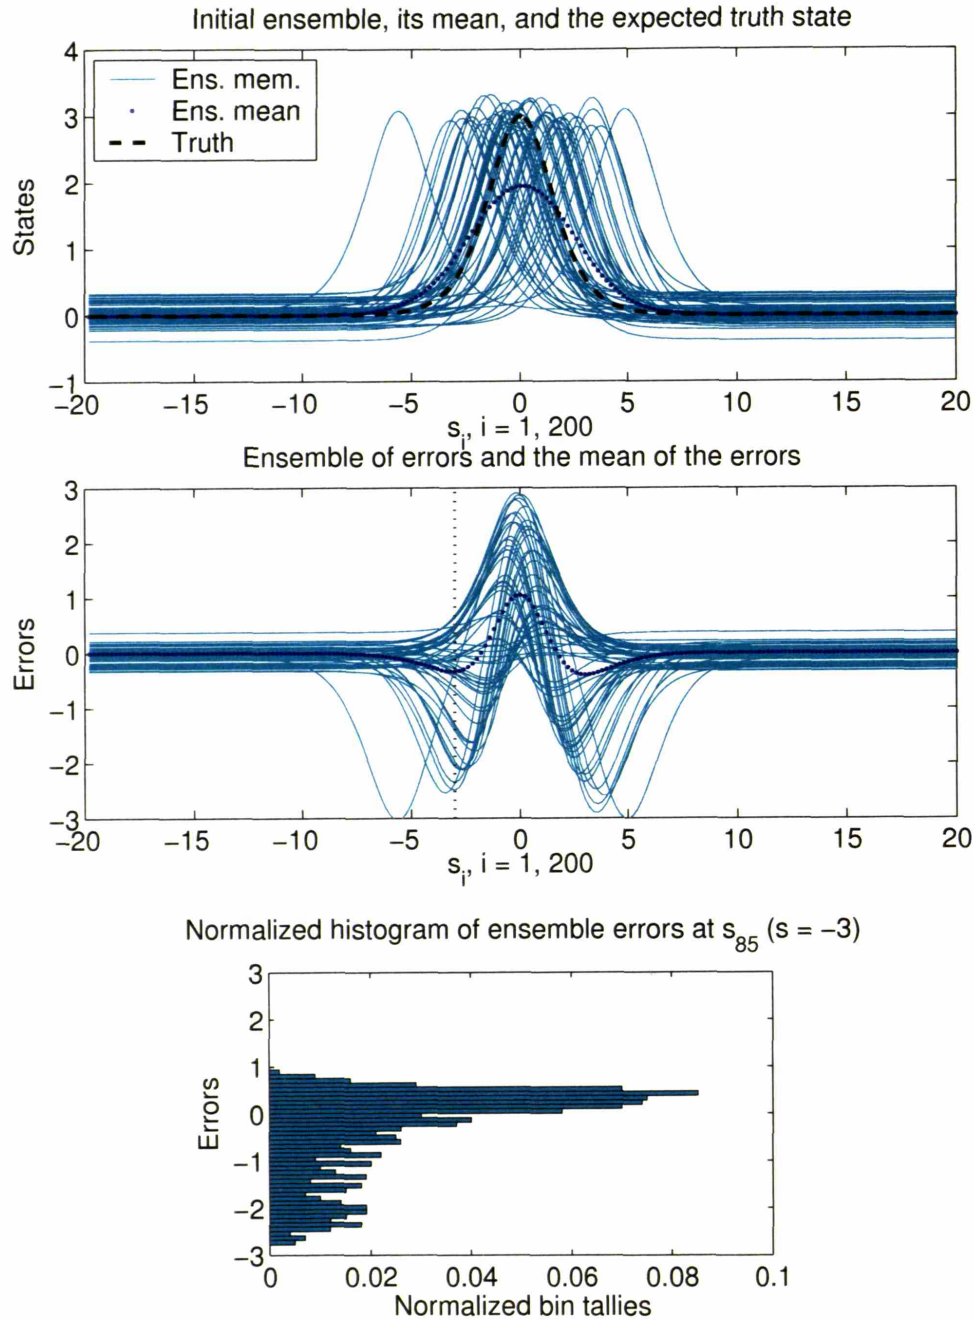


Figure 5-3: The top panel shows a 50 member ensemble generated from the mixed error model described in equation (5.22). The blue dotted line is the ensemble mean, and the black dashed line is the base profile. The middle panel shows the resulting ensemble of additive errors when the ensemble in the top panel is subtracted from the black dashed base profile. The bottom panel shows a normalized histogram of the total additive errors at  $s_{85}$ , (i.e.,  $s = -3.0$ ) compiled from a similarly constructed 1000 member ensemble.

to calculate and visualize. Alternatively, one can simply generate an ensemble of synthetic errors from a strictly Gaussian PDF having the exact first two moments as the ensemble from the mixed error model, and visually inspect the two populations for differences. Doing so, one finds the synthetic total errors (not shown) clearly lack the striking coherence of the mixed error model’s errors (this is shown in the middle panel of figure 5-3). This indicates that the total error’s character depends heavily on non-Gaussian higher moments. A third way to confirm the non-normality was discussed in section 4.1.1, and is shown in the bottom panel of figure 5-3, namely to construct a normalized histogram of the ensemble of total errors for one state element ( $s_{85} = -3.0$  is the element chosen in the figure). Its shape clearly differs from that of a Gaussian.

Also, one can confirm the linearity condition developed in section 5.2 using the same approach. If an ensemble is instead generated with  $\sigma_D = 0.2$ , then the total errors are found to be very well represented by a Gaussian (not shown).

### **Estimation by an additive error model**

It is clear that the total errors for the ensemble in the top panel of figure 5-3 are non-Gaussian, as are they for an ensemble from the mixed error model with additive errors as well. Applying a standard EnKF is not expected to work well. To confirm this, a similarly constructed 50 member ensemble is used as an initial ensemble to be updated. Even though the members were generated by random number realizations, they can be thought of as forecasts from the last time data assimilation was performed. It is important to stress that the discoverer does not know that each of these ensemble members contains a soliton of the specific functional form specified in equation (5.21), the ensemble to the discoverer is only output from a model that is numerically integrating the dynamics because there are presumably no analytical solutions.

Assuming a perfect model, a 51st ensemble member is generated to be truth, meaning the expected value of truth is not used as it is above. The resulting analyses depend crucially on the observations taken and their uncertainty. There are two types of observations considered here, observations of the state itself taken at fixed, known locations, termed “station observations”, and observations of feature position within the state estimate, termed “position observations”.

#### *Station observations*

The example here uses three station observations ( $m = 3$ ), one to either side of the feature, at  $s_{40}$  and  $s_{180}$ , and one within the region of the soliton at  $s_{100}$ . The observations

are equal to the value of truth at those locations plus a random realization of the prescribed observational uncertainty,  $\sigma_{\text{stat}}$ , which here is set equal to the additive uncertainty in the ensemble generation above,  $\sigma_A$ . The observations are assumed uncorrelated, meaning  $\mathbf{R}_{\text{stat}} = \sigma_{\text{stat}}^2 \cdot \mathbf{I}$ , where  $\mathbf{I}$  is the identity matrix. The observation operator is simply an  $m$ -by- $n$  matrix filled with zeros except for a single one in each row, positioned at the indices of the stations. The top panel of figure 5-4 shows the initial ensemble, its mean, truth, and the station observations.

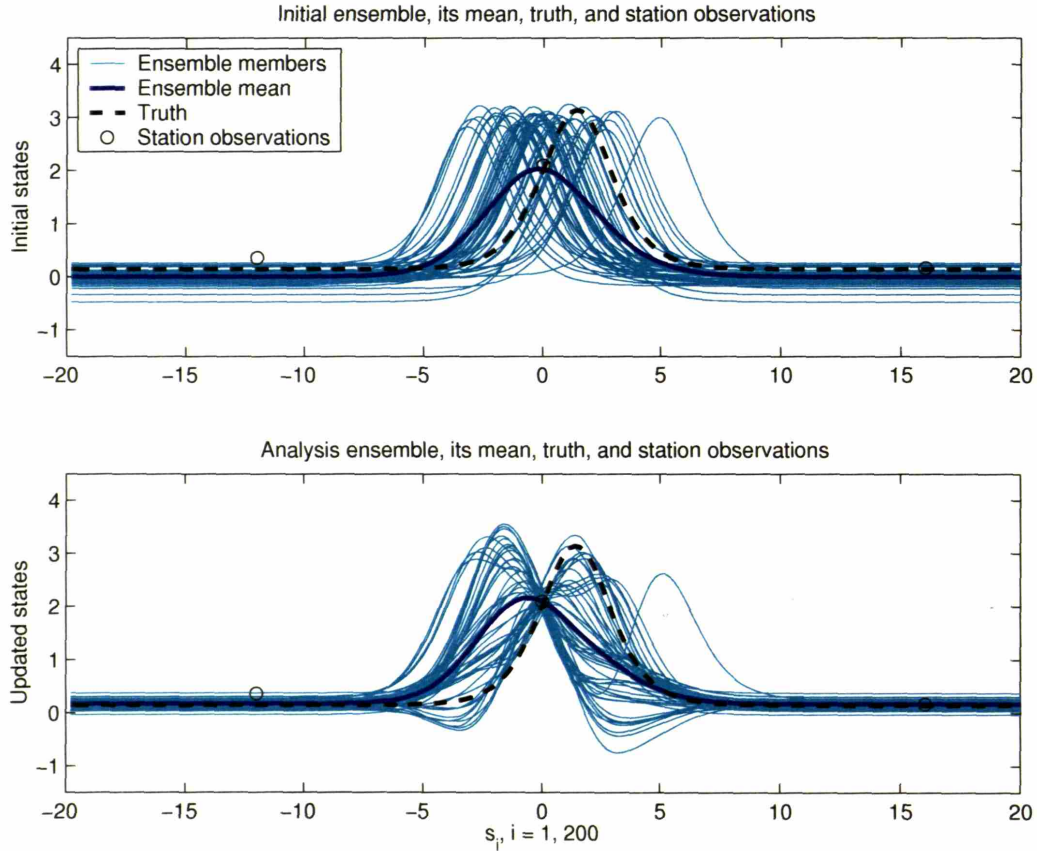


Figure 5-4: EnKF with Station Obs. The top shows the initial 50 member ensemble, the truth profile (simply a 51st generated ensemble member), the initial ensemble mean, and three station observations available. The bottom shows the analysis ensemble from applying an EnKF, its updated mean, truth, and the station observations.

The bottom panel of figure 5-4 shows the analysis ensemble using an EnKF with the station observations. As one central station observation is not a lot of information, there is an ambiguity as to which side of the soliton is being observed, and this is apparent in the analysis. Note that initial members not already consistent with the observation no longer have soliton profiles. This is undesirable if one wants an ensemble of states that

could be random draws from the same PDF from which truth was drawn. If one is only interested in having an ensemble mean with small rms errors, then this analysis ensemble is desirable since it approximates the minimum variance estimate. However, note that the ensemble mean does not look exactly like truth, for example, its amplitude is less and its scale is broader. The ensemble mean is not modified much from its initial state in the top panel because the innovations for the mean are very small. Being designed to update the mean, the EnKF does not change it much at all, yet it does noticeably change the individual members. Note that all ensemble members have been updated so that they pass very close to the central station observation. This is because the observation has  $\sigma_{\text{stat}} = \pm 0.15$  whereas the variances approximated from the 50 member ensemble within the region of the feature give standard errors around  $\pm 0.9$ . Hence, as expected, the non-Gaussian nature of the total error leads to a poor update by a filter that assumes Gaussian errors. Clearly, the covariance of the ensemble alone is not enough to inform the update process that the analyzed ensemble members ought to contain solitons, even though every single member contained one beforehand. It should be mentioned that in a dense observation regime, these problems mostly vanish with the analyzed members simply connecting the station observations. However, in these cases, unless every gridpoint is observed, the ensemble members are not always smooth between station observations.

It is apparent that the filter worked effectively out in the tails of the soliton where the spread about the mean has been reduced and all ensemble members have retained their initial character. The alignment perturbations have a negligible effect on the solitons' tails, leaving only the  $\beta_j$  to contribute to the spread. Additionally, the two side station observations have consistently updated the level of the entire tails, thanks to the covariance information showing that all points outside of the region of the feature covary identically. This comes as no surprise as the  $\beta_j$  are additive Gaussian errors, just what the EnKF was designed to correct.

### *Including a position observation*

The next example adds a position observation to the station observations from the preceding example. An observational error in position of  $\pm 4$  gridpoints, or  $\sigma_{\text{pos}} = 0.8$ , is assumed. In order to compare the ensemble's uncertainty of this, the aforementioned technique of finding  $\mathbf{P}^f \mathbf{H}^T$  and  $\mathbf{H} \mathbf{P}^f \mathbf{H}^T$  directly from the nonlinearly observed ensemble is necessary. The observation operator is thus a nonlinear operator which elicits the value of each ensemble member at the station locations and evaluates the function *arg max* on each member. As the ensemble members are discretized estimates,  $s_{\text{max}}$  will be a quantized

estimate of position, picking out the gridpoint where the maximum occurs.

The top panel of figure 5-5 shows the initial ensemble, its mean, truth, and the various observations. The bottom panel shows the updated ensemble from applying an EnKF. It is clear that some improvement has been made over the analysis in the previous section, both for the mean and the individual members. The ambiguity of which side of the soliton is being observed has been removed, and the variance about the analysis mean has been further reduced. However, many individual members have still been updated into states with features differing from solitons. Again, this is no real surprise.

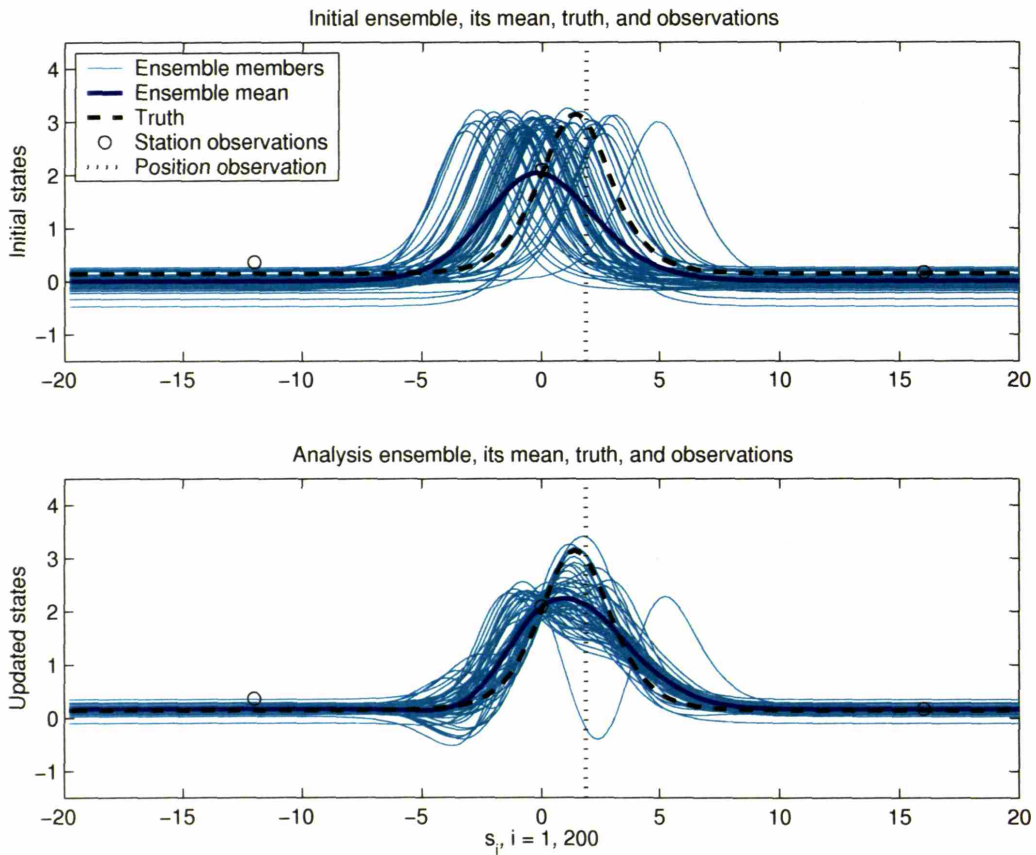


Figure 5-5: EnKF Including a Position Ob. The top shows the initial 50 member ensemble, the truth profile (simply a 51st generated ensemble member), the initial ensemble mean, the station observations, and the position observation available. The bottom shows the EnKF update from using these observations.

## Estimation by an alignment error model

Next we examine how the estimation process is changed by assuming a different error model, a simple alignment error model. This chapter does not address the best way to accomplish alignment error corrections, hence the estimation process used here is extremely basic: it allows for one fixed constant shift, it assumes periodic boundaries, and it only considers position observations. Alignment-based estimation schemes can use station observations (such a scheme is described in the next chapter), but these techniques are beyond the scope of the current example. As the most general form of an alignment error model allows the alignment errors to vary with location themselves, assuming a constant shift in this one-dimensional example here can be thought of as imposing the constraint that the alignment errors be non-divergent.

If one assumes a constant shift alignment error model up front, then the estimation process is redefined as follows:

$$\mathbf{x}^a(s_i) = \mathbf{x}^f(s_i + \Delta s), \quad (5.24)$$

meaning one obtains the analysis by finding the best estimate of position error,  $\Delta s$ , and simply shifting  $\mathbf{x}^f(s_i)$  by that amount. As mentioned above in section 5.3.3, this transforms the problem from the traditional view of updating the prior estimate to one of updating the basis over which the estimate is represented. To clarify this, designate  $r$  as a fixed coordinate system, and now allow the estimate to be represented over  $s$  which is itself a function of  $r$ . So while  $\mathbf{x}^t = \mathbf{x}^t(r_i)$ , the prior estimate is properly written  $\mathbf{x}^f = \mathbf{x}^f(s(r_i))$ , or even further as  $\mathbf{x}^f = \mathbf{x}^f(s^f(r_i))$  to emphasize that the estimation process is now actually:

$$\mathbf{x}^a(s^a) \equiv \mathbf{x}^f(s^a) \quad (5.25)$$

$$s^a(r_i) = s^f(r_i) + \Delta s, \quad (5.26)$$

where  $\Delta s$  is the best estimate of the position error, a random variable.  $\Delta s$  is now the desired analysis increment and can be found from a method like the Kalman filter.

Taking an ensemble approach and using the same position observation as used above, one can apply a standard EnKF to this essentially scalar assimilation problem. The error covariance of the initial estimate is simply the error variance,  $\sigma_s^2$ , of the positions of the ensemble maxima,  $s_{max,j}^f$ . Hence,  $\mathbf{K} = \sigma_s^2(\sigma_s^2 + \sigma_{pos}^2)^{-1}$ . This gain matrix (here, just a scalar) and perturbed position observations give  $s_{max,j}^a$ , the updates of  $s_{max,j}^f$ . These in turn define the various analysis increments,  $\Delta s_j$ , which define the analysis bases,  $s_j^a(r_i)$ , which finally

define the analysis ensemble,  $\mathbf{z}_j^a(s_j^a(r_i)) = \mathbf{z}_j^f(s_j^a(r_i))$ . This is illustrated in figure 5-6.

The top panel of figure 5-6 shows the heart of this estimation problem, namely the shifting of the basis in an effort to better align the ensemble’s features with truth’s. The top line of dots shows the initial ensemble of soliton positions as found from the ensemble shown in the top panel of figure 5-4 with the  $\circ$  denoting its mean. The second line shows the position observation with a horizontal line spanning  $\pm\sigma_{\text{pos}}$  about the observation. The bottom line of dots shows the analysis ensemble of feature positions, again with the  $\circ$  denoting the mean. The middle panel shows the updated ensemble of model states, with each member having been shifted by the amount indicated in the top panel. Note that all ensemble members still have a soliton shape. This shape preservation is built into the error model assumptions. Note also that no correction has been made to any additive state errors (from the  $\beta_j$ ) that exist in the initial ensemble, because an alignment error model alone does not allow for additive errors. Hence, by transforming the estimation process to one of the basis rather than the state, the uncertainty has been reduced in the initial ensemble and each individual member remains a possible random draw from the same PDF that produced truth.

To be able to strictly compare the alignment error model’s results to those from an additive error model, the bottom panel of figure 5-6 shows the analysis ensemble from using the EnKF as in the previous subsection considering only the position observation. Very few of the updated members look similar to solitons. This can be understood by looking at the analysis increments added to each member. As explained in section 5.3.3, the increments are simply scaled versions of the Kalman “gain vector” (because  $m = 1$ ), which closely resembles the spatial derivative of the mean (not shown). One can see how addition of the gain vector to the ensemble mean could shift it quasi-coherently towards the observed position. However, as it is constructed to move the mean only, its structure is not such to be able to move the individual ensemble members effectively. In fact, it is apparent that it fails rather miserably in several cases.

Clearly, neither a purely additive error model nor a pure alignment error model is sufficient for delivering the corrections we suspect to be possible given the initial ensemble and the four observations. The next section explores estimation based on assuming the (proper) mixed error model.

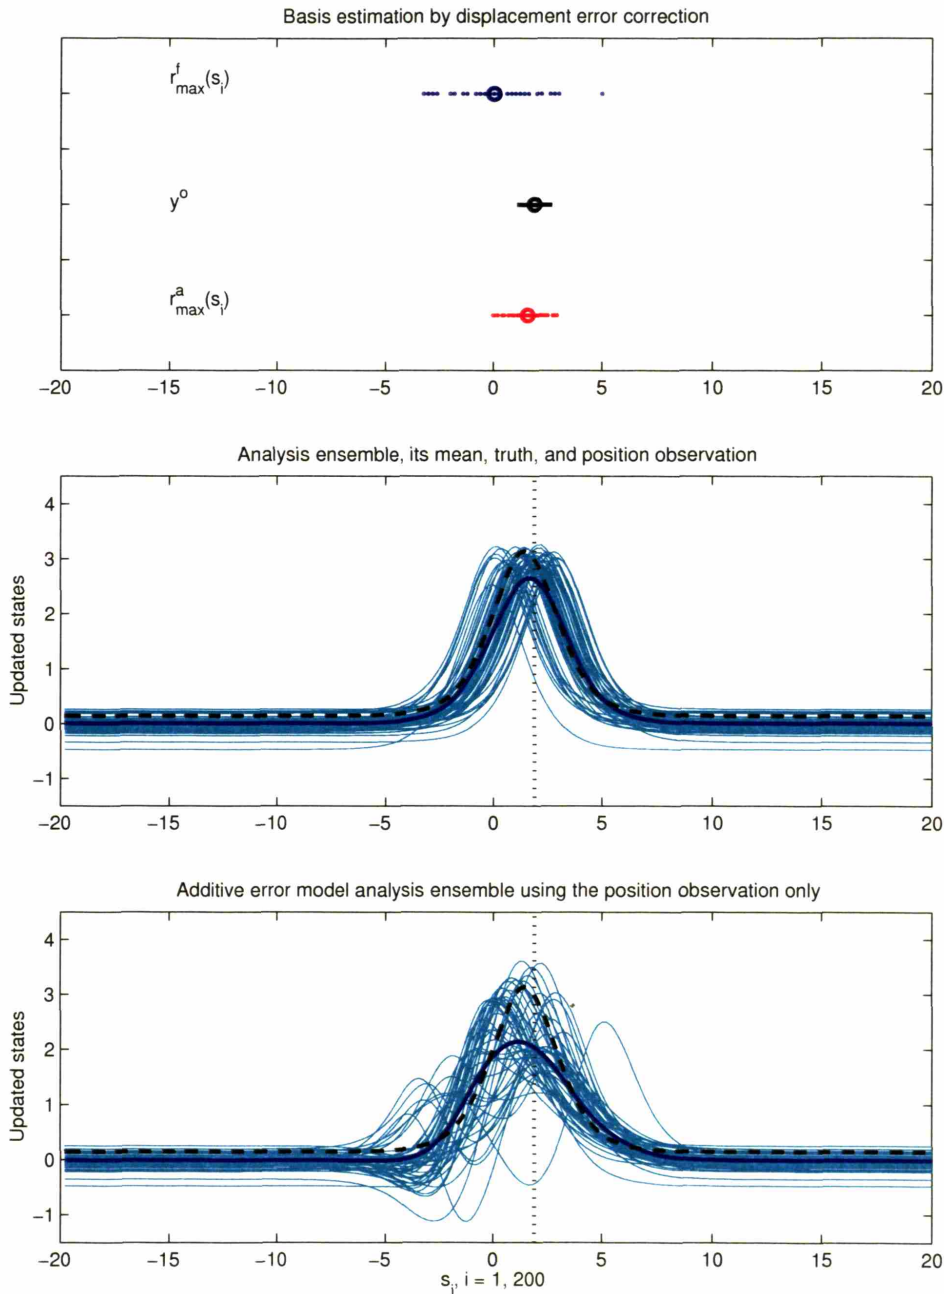


Figure 5-6: Alignment Error Model EnKF. The top shows the mechanics of basis estimation based on the initial 50 member ensemble and the position observation: the top line shows the ensemble of positions of state maxima with its mean denoted by the  $\circ$ , the middle line shows the position observation with its uncertainty denoted by the horizontal bar, and the bottom line shows the EnKF update of the initial ensemble. The middle panel shows the updated ensemble states based on the updated bases in the top. The bottom panel shows the update an additive error model EnKF gives using only the position observation.

## 5.4 A two-step filter approach

The previous two sections have introduced a mixed error model example and shown how assuming an improper error model can harm the estimation process; this section shows the benefits of assuming a more natural error model. The analysis ensembles in the examples of the previous section indicate that consideration of the proper error model naturally leads to a two-step approach. The total additive errors of the initial ensemble are non-Gaussian due to the relatively large position errors (see top panel of figure 5-4), but the total additive errors of the analysis ensemble from the alignment error model filter are much more Gaussian (see third panel of figure 5-6). Hence, by using an alignment error model estimation scheme first, subsequent application of an additive error model estimation scheme will work successfully if the first step has corrected the position errors such that the total additive errors are in their Gaussian range. The order of these two steps is clearly important because application of any data assimilation scheme, ensemble-based or not, stemming from the Kalman filter ultimately assumes the errors are at least approximately Gaussian. Hence, the cause of the non-Gaussianity must be addressed first — even a simple Gaussian position error can lead to non-Gaussian additive errors. If the non-Gaussianity is thought to stem from the mixed error model under consideration, then this two-step approach is a way to retain usage of traditional filtering methods. If its assumptions are met and used in an ensemble context, this two-step approach will achieve both of the goals of ensemble filtering mentioned in the previous section: to have an ensemble mean with small errors and to have each ensemble member be statistically indistinguishable from truth.

Figure 5-7 shows the application of this two-step approach to the examples considered in section 5.3.4. Beginning from the same initial ensemble as the examples above (see the top panel of figure 5-4), one first uses the alignment error model EnKF as shown in figure 5-5. The top panel of figure 5-7 shows the analysis ensemble from this first step, its mean, truth, and the station observations (i.e., this shows the same information as the middle panel of figure 5-6 with the addition of station observations). The bottom panel shows the further updated ensemble from using the station observations and a standard additive error model EnKF (*cf.*, figure 5-4). The final updated ensemble looks reasonable, but there are still some members that have lost their initial soliton shapes, though not to the same degree as shown in the bottom panel of figure 5-4. The simple linearity condition developed in section 2 could have been used to deduce that the intervening analysis ensemble in the top panel of figure 5-7 does not quite have Gaussian total errors, but regardless, the discoverer must decide if the ensemble updates are satisfactory.

As explained in sections 2.3.3 and 5.3.2, the EnKF favors the ensemble mean at the

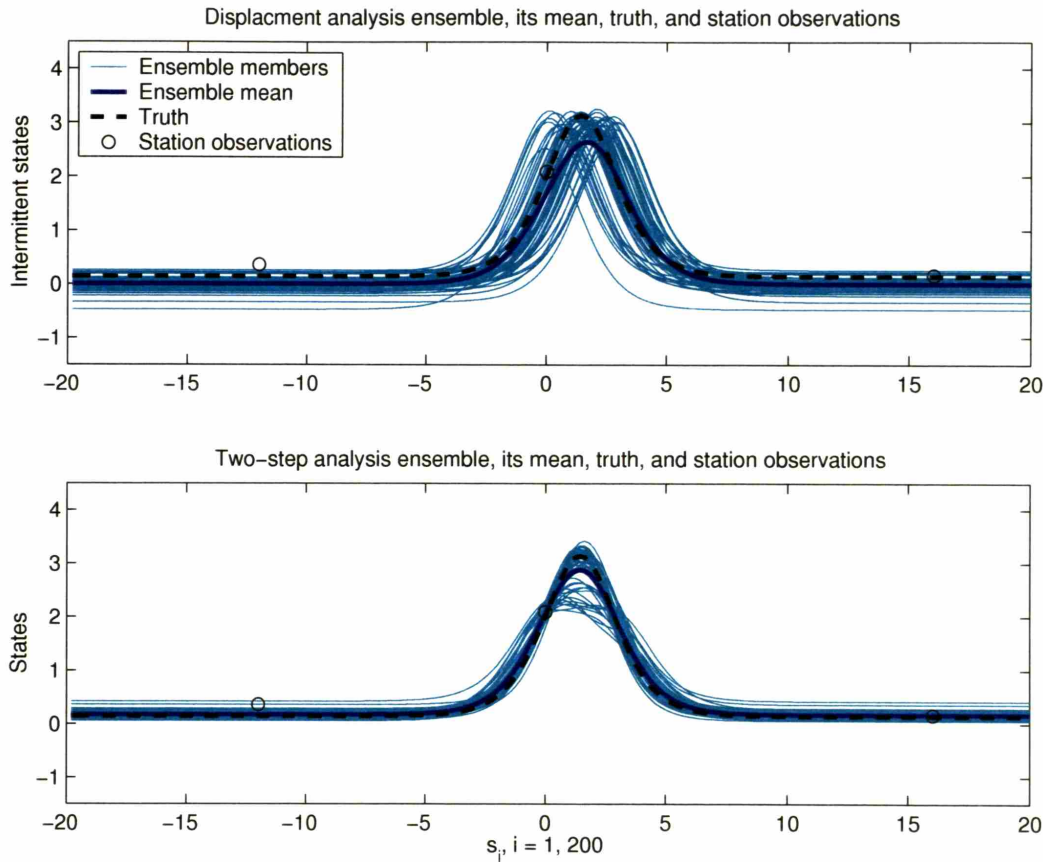


Figure 5-7: Two-Step EnKF. The top shows the analyzed ensemble from having assumed an alignment error model and using the position observation (from the bottom panel of figure 5-6), the ensemble mean, truth, and the same three station observations as used before. The bottom shows the updated ensemble from further applying an additive error model EnKF.

expense of the individual members. If the loss of soliton shape in the individual members is deemed undesirable, then strictly, another estimation method must be found as the errors in this estimation problem, at least based on the mixed error model assumed, are simply not correctable by Kalman filter-based methods. However, if one would like to maintain use of these methods and is willing to favor the individual members over obtaining an ensemble mean with the desired analysis uncertainty, then an approximation to the two-step approach can be used. The essence of the proper approach is to correct the position errors before trying to correct the additive errors. This spirit can be captured by artificially reducing or eliminating the position errors before correcting the additive errors. Of course, having position errors is an important part of the uncertainty of the ensemble estimate, so the errors should be reintroduced after the additive errors are corrected. This artificial feature

co-location across the ensemble has the effect of imprinting the feature into the ensemble mean, thus creating a stencil of sort for the application of an additive error model based estimation scheme. The method of achieving this feature co-location is somewhat arbitrary since no error statistics exist to guide the process; one could even use morphing methods from image processing, such as those performed by Alexander et al. (1998), as long as the positioning is consistently unraveled after the additive errors have been corrected.

In the simple examples studied here, image morphing techniques can be avoided as the soliton profiles can just be shifted left and right until their features are aligned. An important assumption that has been made here is that alignment errors are uncorrelated with, and therefore separable from, additive errors. A better alignment error model estimation scheme than the one used here would not assume this and would utilize all available observations, including station observations, to correct the position errors. In the examples above, however, the alignment error model filter used is only equipped to utilize position observations. Therefore, as the example below demonstrates, when reintroducing position uncertainty back into the ensemble, inconsistencies can arise due to the separability assumption. That is to say, additive errors corrected in the co-located ensemble may not appear corrected in the un-transformed ensemble.

Figure 5-8 shows the application of this co-location approximation to the two-step procedure. After one has performed alignment error assimilation (see top panel of figure 5-7), the top panel shows how one might choose to artificially co-locate the features in each ensemble member at the analysis mean position,  $\bar{s}_{max}^a$ . The middle panel shows the analysis from applying a standard EnKF to the co-located ensemble using the station observations. The bottom panel shows the final analysis ensemble once the artificial co-location procedure has been undone. Note the inconsistency of ensemble spread in the vicinity of the middle station observation. Each analyzed ensemble member in the proper two-step approach is consistent with the specified uncertainties of the problem; consequently, some have lost their initial feature shape. By forcing the two-step approach with the co-location approximation, the estimation process has begun to favor the individual ensemble members at the expense of consistently satisfying all uncertainties, that is, at the expense of minimizing uncertainty in the mean. The analysis ensemble in the bottom panel has non-Gaussian errors, but this need not cause concern, as the initial ensemble had significantly non-Gaussian errors; the two-step approximation forcibly treats the non-Gaussian total errors according to the assumed error model.

Having examined the mechanics of the two-step approach as applied to contrived, static examples, the next section examines its behavior in a dynamic setting by way of OSSEs.

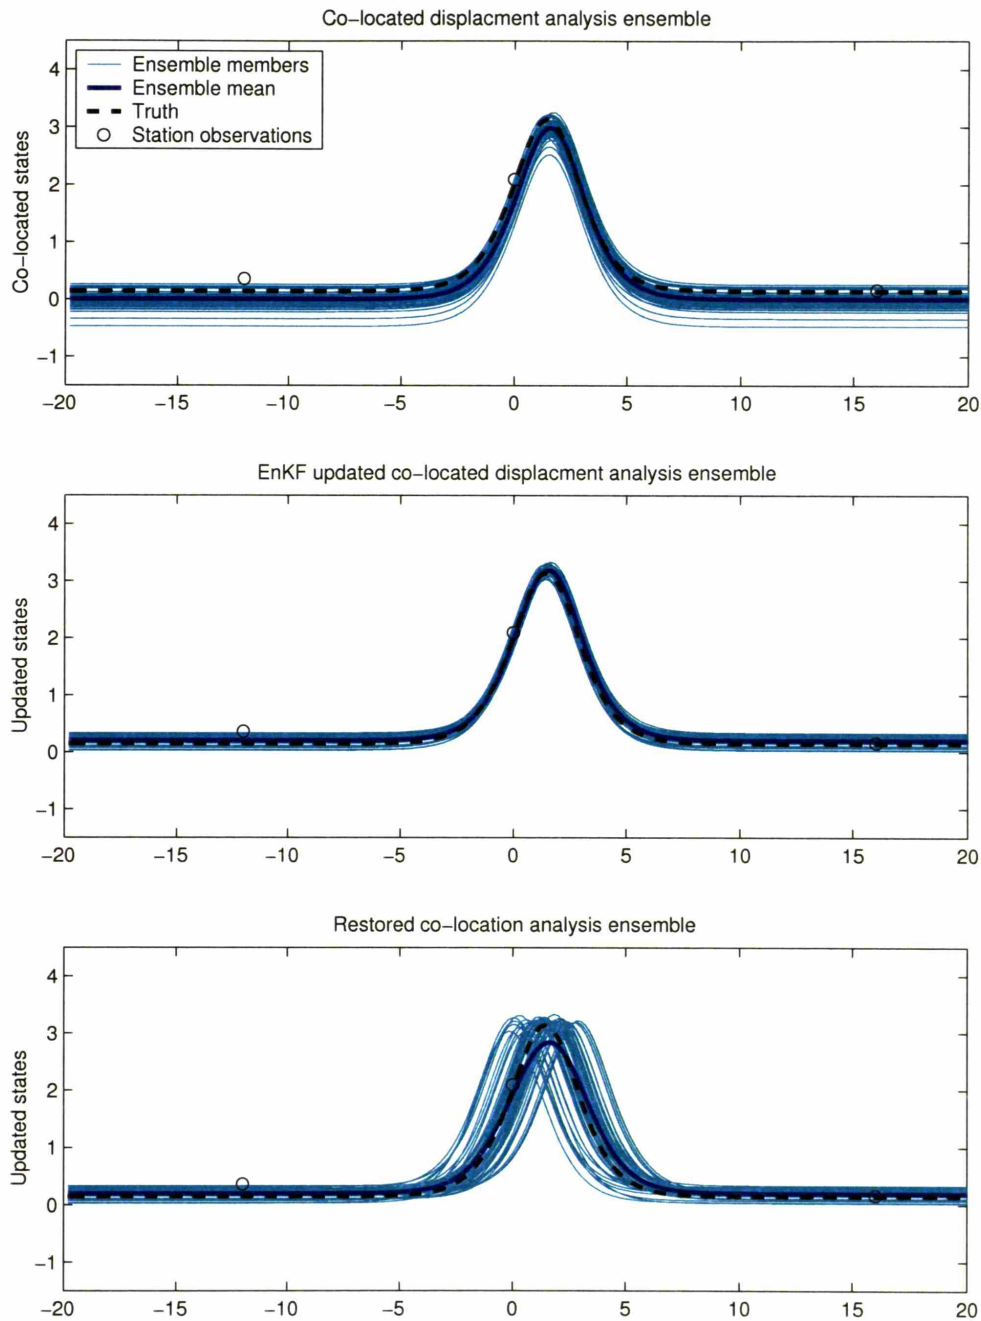


Figure 5-8: Co-location Two-Step Approach. The top shows the ensemble with every member artificially shifted so that their feature positions are all co-located with the ensemble mean location as determined by the alignment error model EnKF (*viz.*, the middle panel in figure 5-6). The middle panel shows the update of co-located ensemble by an additive error model EnKF using the three station observations. The bottom panel shows the updated ensemble in the middle panel with each member restored according to its initial feature position.

## 5.5 Observation system simulation experiments

In this section, OSSEs are presented using the KdV equation and various estimation schemes. The update examples in previous sections are static; the examples presented here are dynamic. Including the forecast step shows how problems can compound through the assimilation process. Also, since solitons have phase speeds proportional to their amplitudes, if the initial solitons differ in amplitude, the forecast step will allow systematic feature dispersion to occur. The pseudo-spectral, nonlinear, implicit time-stepping scheme of Li and Sattinger (1998) is used to numerically integrate the KdV equation.

### 5.5.1 Experimental set-up

To impart the ability for features to disperse in position, the ensembles considered here include calibration errors, that is, multiplicative errors, as well. Estimation still proceeds assuming the mixed error model of equation (5.6) even though the initial ensembles now actually follow the error model in equation (5.5). To simplify the problem, as with the additive errors, the vector of multiplicative errors is assumed to perfectly covary throughout the domain, meaning only one random number realization per ensemble member is necessary,  $\gamma_j$  from the PDF  $N(0, \sigma_M^2)$ . Ensemble generation thus follows  $\mathbf{z}_j^f = (1 + \gamma_j)\mathbf{1} \circ \mathbf{x}^b(s_i + \alpha_j) + \beta_j\mathbf{1}$ ,  $j = 1, N$ . By uniformly amplifying each soliton, a risk is run of significantly altering the base soliton profile, possibly leading to additional, unplanned, dispersive wave forms. For this reason and so that the assumed mixed error model is not completely inaccurate, the multiplicative errors are kept relatively small throughout.

For experimental design, the creator chooses to have a domain size of 40 centered at 20 using 512 gridpoints (because of the pseudo-spectral time-stepping routine). The constant  $A$  in equation (5.21) is chosen to be 4, making the features even more compact than in the static examples above (half-width of 1.8). The time step is set at  $\Delta t = 0.005$ . Observations available include 8 fixed, equally spaced station observations, each with assumed uncorrelated uncertainties, and a position observation, again assumed uncorrelated to the station observations. For initial ensemble generation, the examples here use:  $\sigma_D = 5$  gridpoints, or  $= 0.39$ , to begin position discrepancies in the nearly Gaussian total error range, natural dispersion over the integration will take the ensemble out of this range;  $\sigma_A = 0.6$ , again equal to 5% of the total amplitude; and  $\sigma_M = 0.05$  to allow for ensemble dispersion but not strongly deform ensemble members from the analytical form of soliton solutions. An additionally generated ensemble member is designated truth, thereby allowing truth to differ slightly from the exact analytical form in equation (5.21). The number of members

generated is again chosen to be  $N = 50$ .

No attempt is made to link chosen time and spatial scales to anything physical; however, the time between observations,  $\tau_{obs}$ , and the observational uncertainties,  $\sigma_{stat}$  and  $\sigma_{pos}$ , provide controls over the error growth, and these controls are chosen to roughly mimic NWP applications of interest, that is, to allow error growth to become slightly nonlinear (i.e., PDFs to become slightly non-Gaussian). As the unforced/undamped KdV equation is not chaotic per se, the nonlinear error growth is mainly due to growing position errors. Therefore, when the position errors come to violate the linearity condition developed above, the two-step EnKF should work well where a purely additive error model EnKF works poorly.

### 5.5.2 The failure of the tangent linear hypothesis

Before considering the ensemble updates, we first take this opportunity to demonstrate an aspect of position errors that can cause methods that rely on the dynamics' tangent linear model (TLM, see equation (2.34)) to fail. An example of a commonly used scheme that implements the TLM is the EKF (referred to above in 5.3.2 and described in section 2.3.3). The EKF is a deterministic DA scheme (rather than a probabilistic scheme like the EnKF), and so only evolves a single state estimate,  $\mathbf{x}$ , and its covariance matrix  $\mathbf{P}$ . The potential problem with using the TLM in the face of features is that the TLM ensures linear error growth, whereas *additive* errors can grow quite nonlinearly in a short amount of time. In particular, for a sharply enough defined feature, a small position error accrued over a short time can lead to a large additive error, and this error amplification will not be captured by the TLM. At DA time, this can lead to too confident a prior estimate, which in turn can lead to bad updates and filter divergence.

To demonstrate this, we have coded the TLM of the KdV system so that we can implement the EKF. The EKF evolves the covariance matrix, so we need a sensible covariance matrix with which to begin. For this purpose, we simply used a very large ensemble generated as discussed in section 5.5.1 so that the initial estimate of  $\mathbf{P}$  approaches the correct covariance matrix. Figure 5-9 shows one such incidence when the TLM failed. For this case we used:  $\sigma_{stat} = 0.6$  (no position observations were used as their corresponding observation operator cannot be sensibly linearized: see section 4.2.4) and  $\tau_{obs} = 250\Delta t$ . The top panel shows the initial truth state (black) and the initial EKF estimate (blue). Both are random realizations from the described ensemble generation method in section 5.5.1. The second panel shows the short term integration to the next observation time (the dotted line shows how far the solitons have propagated over the integration time). Note that the EKF

estimate has lagged slightly behind due to its slightly smaller amplitude. The third panel shows the station observation locations and values (the black circles) and the resulting EKF analysis (red). It is a particularly bad analysis because the station observation falls right at the peak of the estimate’s soliton, and the TLM has not allowed the estimate’s uncertainty at its peak to change accordingly. To emphasize the danger of bad analyses in a cycling DA environment, the bottom panel shows the ensuing forecast that would result from further integrating the analysis until the next observation time. We note that this is somewhat of a special example because the estimate’s peak happened to be right over the observing station; probably in a more realistic setting an observational quality control algorithm would not allow this sort of analysis to take place. However, we also note that we did not have to go searching for this failure, it was quickly obtained after the EKF was coded and experimentation had begun. This example makes a further case for why ensemble methods are preferable in these types of scenarios since they avoid linearizing the dynamics.

### 5.5.3 Ensemble updates

To confirm sensible behavior of the system within a regime of linear error growth, the first experiments are run with frequent and accurate observations. Using  $\sigma_{\text{stat}} = 0.6$ ,  $\sigma_{\text{pos}} = 0.25$ , and  $\tau_{\text{obs}} = 100\Delta t$ , the standard EnKF, the two-step EnKF, and the co-location approximate two-step EnKF perform essentially identically. This is because the position errors are well-contained within the linear regime. Further comparisons were made applying an EKF in this regime, and, as expected, it works just as well, even though it cannot sensibly use the position observations. No figures are shown for these experiments because they would simply show a not-too-dispersed ensemble tracking truth very well. The same results are achieved using  $\sigma_{\text{pos}} = 0.50$ .

Keeping the observational errors set at  $\sigma_{\text{stat}} = 0.6$  and  $\sigma_{\text{pos}} = 0.25$ , errors are allowed to become non-Gaussian by extending the time between observations to  $500\Delta t$ ’s. The top panel of figure 5-10 shows the initially generated ensemble and truth. Note that the position dispersion is kept small and that the ensemble mean appears very similar to truth. The bottom panel shows the ensemble and truth after they have been integrated  $500\Delta t$ ’s. Because the ensemble members now have slightly differing amplitudes, their phase speeds are slightly different, and hence the initially small position dispersion has increased to as shown. The total errors are outside their Gaussian range. The bottom panel also shows the nine observations taken so that data assimilation may proceed.

Figure 5-11 shows the updates of the integrated ensemble in the bottom panel of figure 5-10 by the traditional EnKF (top panel) and the two-step approach (bottom panel). Note

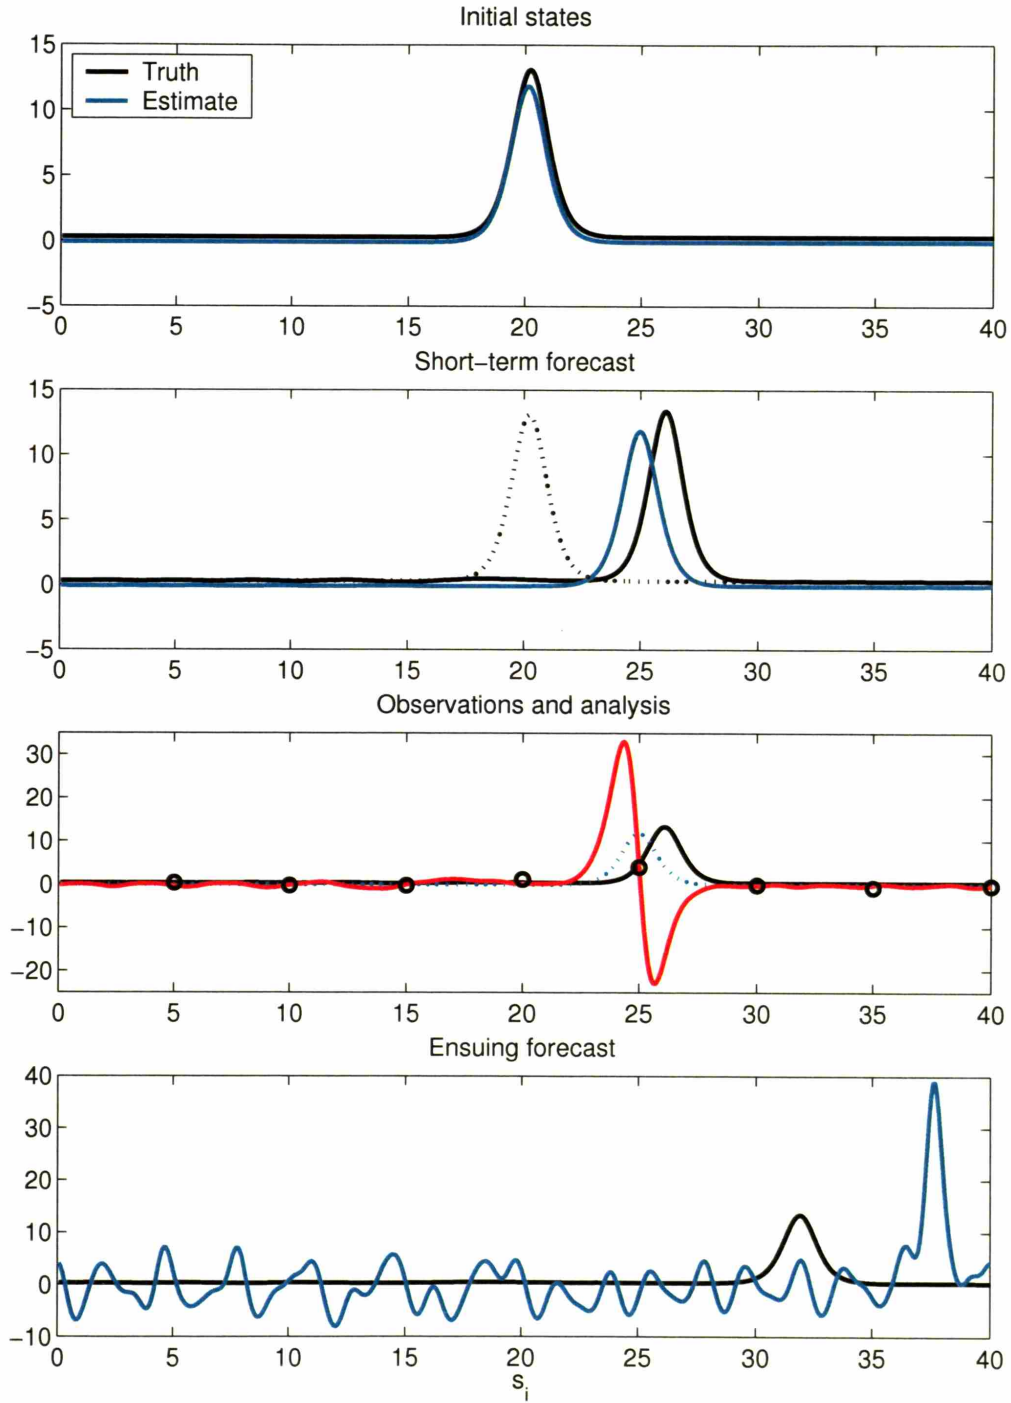


Figure 5-9: EKF Example. The top panel shows the initial truth and estimate states. The second panel shows their short term integrated states. The third panel shows the eight station observations that were made available and the resulting EKF analysis state in red. The bottom panel shows the ensuing time integrations.

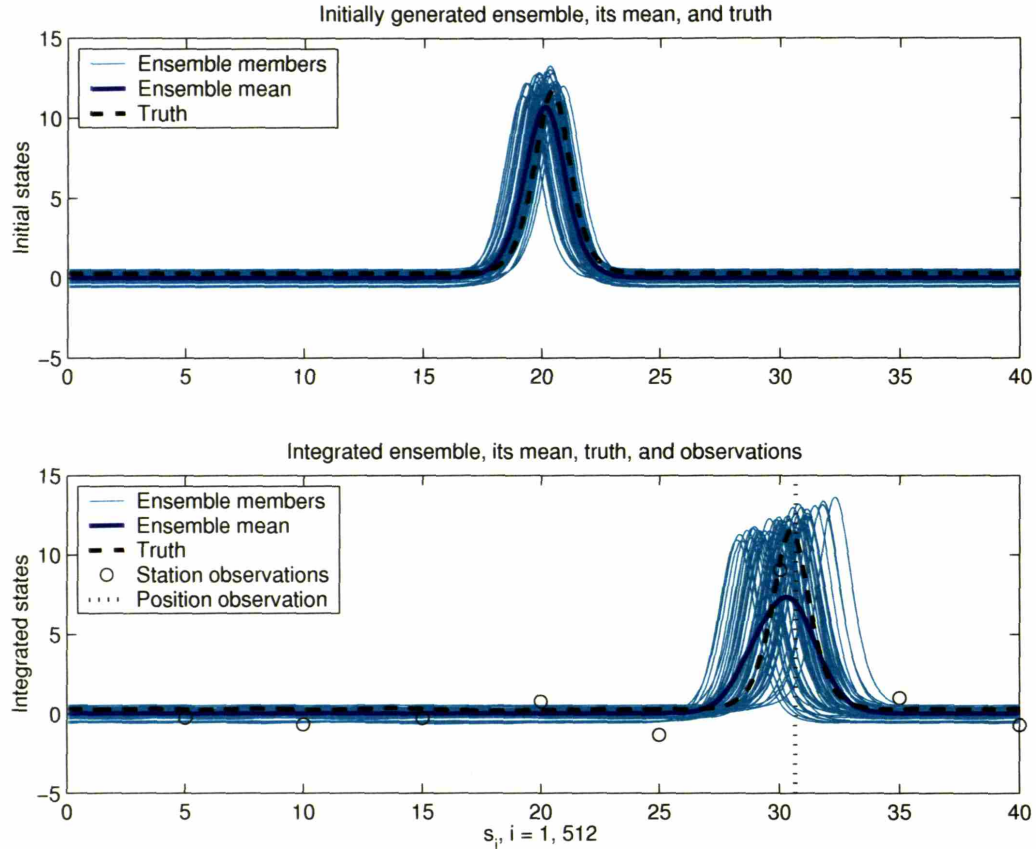


Figure 5-10: The top shows the initially generated ensemble, its mean, and truth. The bottom shows the ensemble from the top panel having been integrated under the KdV dynamics for  $500\Delta t$ 's. The observations that will be used in the update step are also shown.

in the top panel, as expected, that many of the updated members no longer contain soliton forms. The ensemble mean is a reasonable representation of truth, but the individual members are no longer statistically indistinguishable from truth. The ensemble members look much better in the bottom panel, all ensemble members still look like solitons, and it is clear the ensemble mean has small errors. The two-step approach is able to produce this desirable analysis because the position observation is accurate enough to render the total errors Gaussian after application of the alignment error model EnKF.

To see how the updates in the top panel of figure 5-11 may harm the filtering process, the analysis ensembles in figure 5-11 are further integrated to the next time observations become available,  $500\Delta t$ 's later. Figure 5-12 shows the ensuing forecasts for each ensemble. Note the top panel shows spurious waveforms developing in regions away from the soliton. These are the results of unbalanced dispersion and advection of and by the “wiggles” formed by the

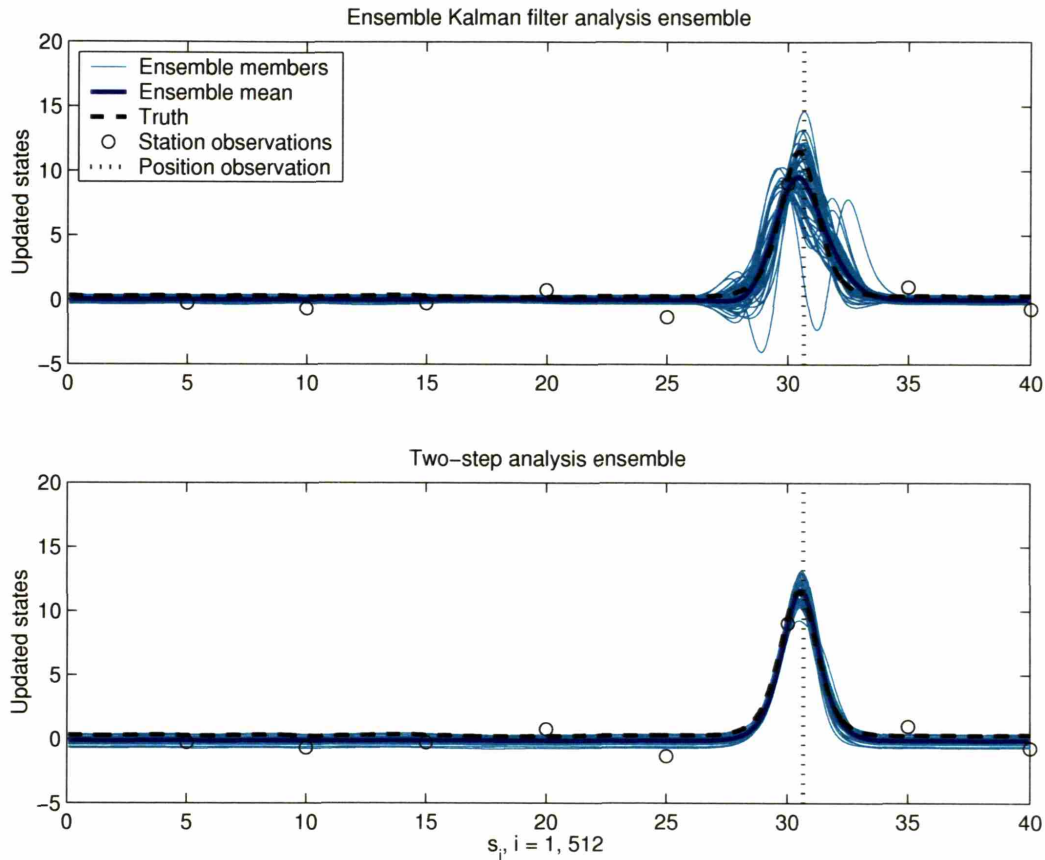


Figure 5-11: The top shows the analysis ensemble from the traditional EnKF. The bottom shows the analysis ensemble from the two-step approach.

data assimilation system. The two-step EnKF ensemble in the bottom panel has evolved almost exactly as it does between the two panels in figure 5-10, a result of maintaining properly shaped features through the assimilation process. Though the traditional EnKF ensemble in the top panel appears to bound truth and give reasonable representations of the state, the effects of non-Gaussian errors and their ensuing nonlinear growth quickly compound.

Continuing the filtering process, the ensemble being updated by the EnKF begins to lose correspondence to truth by its fourth assimilation. By the twelfth assimilation, shown in the top panel of figure 5-13, it has ceased to bear resemblance to the soliton solution of truth. As the state updates continue, the variance about the ensemble mean is reduced until all members essentially collapse onto the mean, at which point the filter has diverged and further observations are rejected. This occurs by the twentieth assimilation; the states are then simply freely-evolving solutions of the KdV equation. The bottom panel of figure 5-13

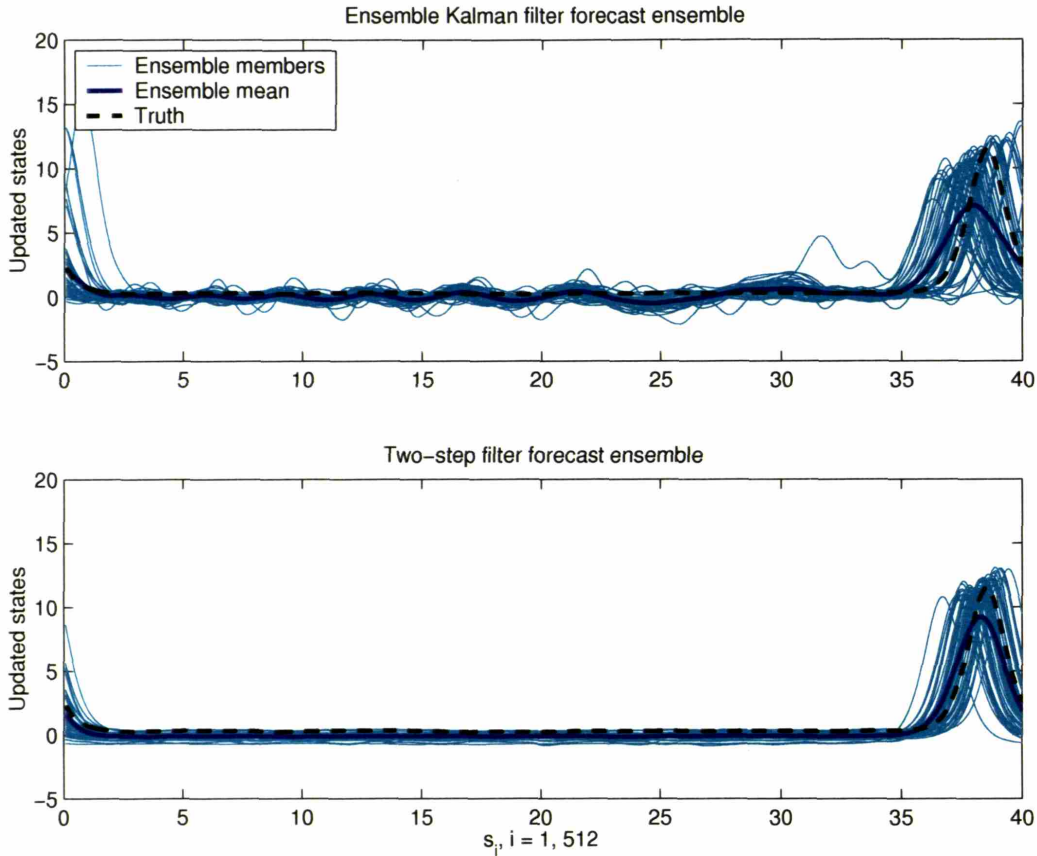


Figure 5-12: These show the forecast ensembles from the different filters' analysis ensembles shown in figure 5-11. The ensembles have been integrated another  $500\Delta t$ 's to the next observation time.

shows the analysis ensemble from the two-step approach. Its ensemble mean is very close to truth, and each ensemble member still contains a well-shaped soliton in approximately the correct location.

Were the position observation less accurate, the two-step approach would not work as successfully because the total errors could still be non-Gaussian after the first step. Figure 5-14 shows the analysis ensembles of the EnKF and the two-step approach after twelve updates using position observations with  $\sigma_{\text{pos}} = 0.50$ . The EnKF has lost correspondence to truth even more quickly than before, and the two-step approach filter is in the process of diverging. Both filters have completely diverged by the twentieth update. The two-step filter performs better than the traditional EnKF and for longer, but it still eventually fails due to the accrued effects of the non-Gaussian errors. The bottom panel shows the analysis ensemble from having used the co-location approximate two-step filter. The approximate two-step filter is not observed to lose correspondence to truth, even well beyond the twenti-




Figure 5-13: These show the analysis ensembles after twelve assimilations for the two different filters.

eth assimilation. However, each of its analysis ensembles is not completely consistent with all information available for its update.

These OSSEs show and confirm that the effects of trying to correct non-Gaussian errors with linear methods can quickly compound and derail the filtering process, though ground can be gained by assuming an alternative error model.

## 5.6 Conclusion

The concept of alternative error models has been suggested as a means to redefine estimation problems with non-Gaussian errors so that familiar, near-optimal methods may still be successfully applied. The specific example of a mixed error model including both alignment errors and additive errors has been examined. Using the specific form of a soliton, an analytical solution to the KdV equation, the total errors of states following the mixed error

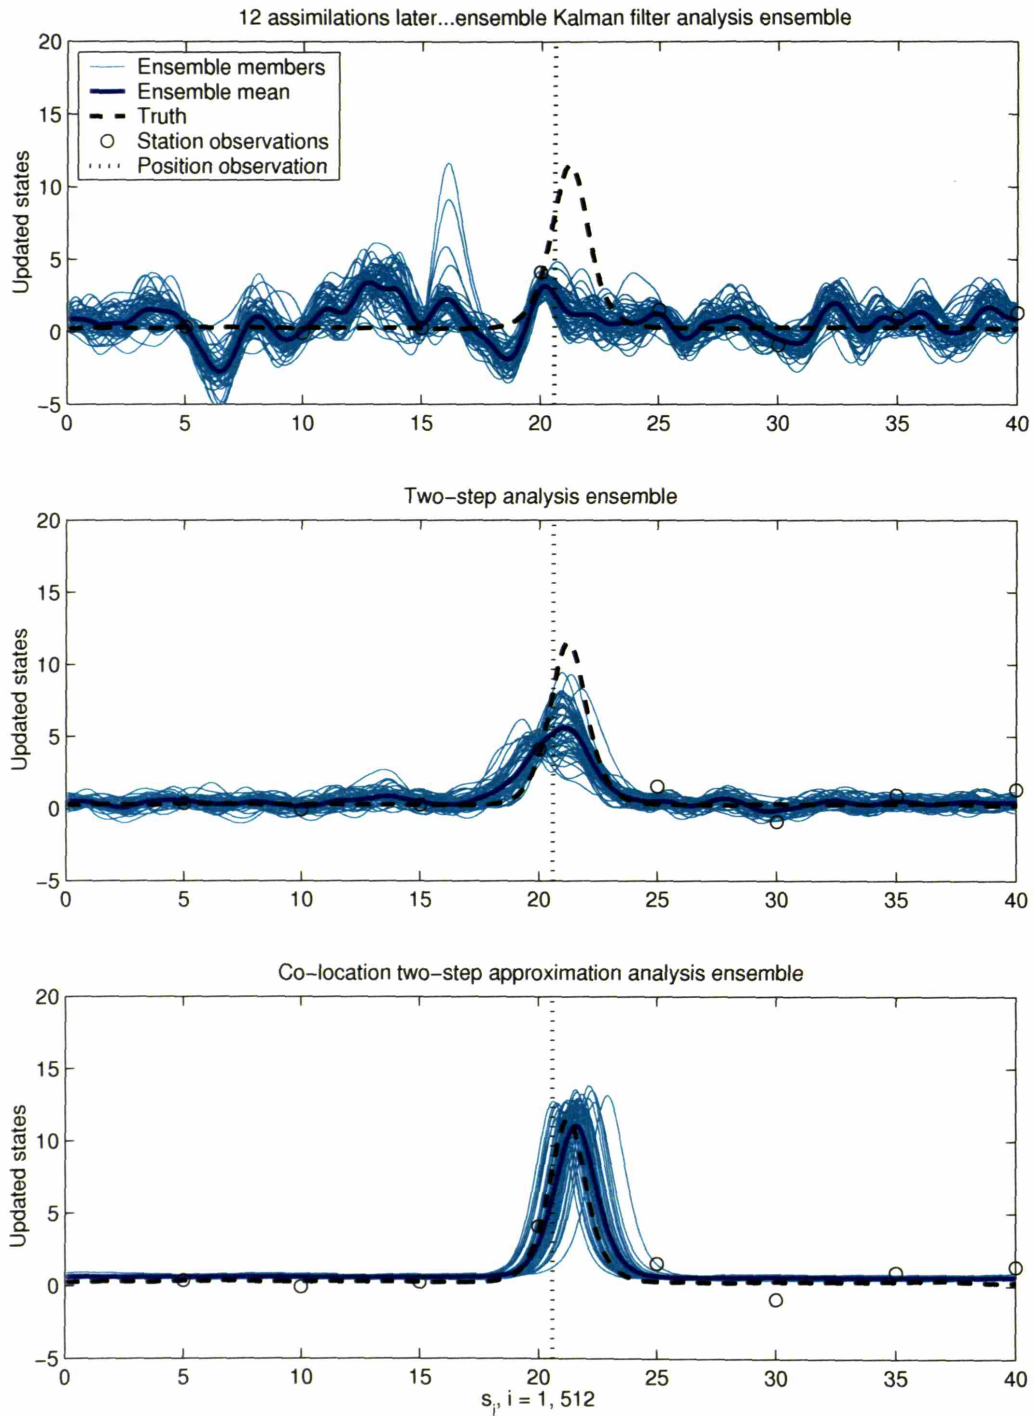


Figure 5-14: The top and middle panels show the same information as figure 5-13, the analysis ensembles after twelve assimilations for the two different filters, only having used an observational error twice as large. The bottom panel shows the comparable analysis ensemble from the co-location approximate two-step approach.

model have been shown to be non-Gaussian, and an ensemble of such states has been shown to be handled poorly by a traditional EnKF, even if position observations are included. Consideration of the mixed error model naturally suggests a two-step approach, where the alignment errors are corrected first followed by application of an estimation scheme to the remaining additive errors. If the first step removes most of the non-Gaussianity, then the second step can proceed successfully, and the final analysis will approximate the desired minimum variance estimate.

Again taking an ensemble approach for the soliton states, this two-step approach has shown a great improvement. For the alignment error correction step, a very simple EnKF scheme has been used that considers only position observations. More sophisticated alignment error estimation schemes of course exist that take advantage of all available observations (as we explore in the next chapter); *such schemes are not the focus of this study, and the simple one implemented here is sufficient for showing the two-step approach's effectiveness, thereby confirming the power of considering alternative error models in estimation.*

If the first step of the two-step approach has not removed the non-Gaussianity of the total errors, the approach will not work in a strict statistical sense. However, a convenient approximation is available based on the principles that make the proper approach successful. The approximation involves an artificial step where the source of non-Gaussianity, that is, the displacement errors, are eliminated before the additive errors are addressed. The alignment errors are re-introduced after the additive errors have been corrected. Analysis ensembles from this approximate method still have non-Gaussian total errors and therefore do not approximate the desired minimum variance estimate sought by linear methods, but then, the non-Gaussianity of the ensemble's initial errors would have led to poor analyses had the approximation not been made. In an ensemble setting, the approximation serves to shift the "badness" of the analysis. Instead of having un-physical updates within the individual ensemble members, one has physically sound individual members with an ensemble mean no longer well-approximating the minimum variance estimate. Clearly, this approximate two-step method is not desirable to have to implement, but it can be useful in cases where the user values realistic ensemble member updates over obtaining the most accurate ensemble mean estimate.

The experiments presented here comprise a very simple demonstration meant only as a proof-of-concept. There are many levels of imaginable "real-world" complications to be added. One simplification used above is that a position observation is used to re-position the feature; as already mentioned, using station information would be more powerful. Further simplifications arise from assuming a constant shift in alignment, for example, all ensem-

ble states include only one feature and each feature is succinctly represented by a single position. One can easily imagine having to satisfy multiple, possibly conflicting, position observations, or trying to reconcile differently shaped features in which case the alignment error model may want to address feature position and its spatial gradient. Another assumption made throughout is that of a perfect model. Were one to include sources of model inadequacy in the examples presented here, one could try compensating for them in the data assimilation by the usual methods, for example, including a model error covariance term (i.e.,  $\mathbf{Q}$ ) in the formulation of the gain matrix (equation (5.13)), performing stochastic ensemble integrations with random samples from a specified model error covariance term, or simply by artificially inflating the ensemble covariance. An alternative, interesting approach would be to try compensating for model inadequacy with a judiciously chosen error model, as was intimated in the hypothetical example of parametric error in the introduction.

Though this chapter has chosen simplicity, both in its examples and its alignment estimation scheme, it has shown that consideration of alternative error models can be useful in estimation problems. The specific example of a mixed error model containing alignment errors is an illuminating one, but it has also been motivated as a useful one based on its being physically appealing. In the next chapter we relax the constant shift assumption and extend the framework to more realistic, two-dimensional problems.

# Chapter 6

## Extensions to multiple features and multiple spatial dimensions

This chapter marks the culmination of the preceding chapters' work. In chapter 2, we grounded and extended Bjerknes's sense that geophysical fluid dynamics is merely a complicated dynamical system (1904), and we discussed how this renders numerical weather prediction (NWP) a probabilistic filtering exercise. A theoretically correct approach to probabilistic filtering exists, namely evolution via the forward Fokker-Planck equation and state estimation via Bayes rule, but practical limitations in both computing and our ability to supply the necessary information prohibit the correct approach's application to all but the most academic of problems. Chapter 2 considers the various simplifications of the correct approach that have been posed, and it ends with descriptions of the state-of-the-art data assimilation (DA) systems, namely ensemble-based Kalman filtering approaches. The original Kalman filter was posed for systems whose dynamics are linear and whose (random) errors are accurately described by Gaussian probability density functions (PDFs). Ensemble-based implementations of the original filter allow extension into nonlinear and non-Gaussian regimes, but only weakly so as the filters are still built upon the original assumptions of linearity and Gaussianity. Chapter 3 explored the limits of linearity and Gaussianity more thoroughly for different popular posings of these state-of-the-art filters, namely stochastic filters and deterministic filters. Taking a hierarchical approach through models of varying complexity, it was found that the filter posings fail differently as nonlinear error growth becomes more pronounced, but the important point to note is that they each fail at some level of nonlinearity and non-Gaussianity, as they must. Also, there are ranges of parameters and observational coverage where the filters may still give accurate analyses even when they are no longer giving probabilistically correct (i.e., reliable) analysis

distributions, that is, the ensemble analysis mean may be close to truth even though truth is statistically *distinguishable* from the individual ensemble members.

While chapter 3 looked at generic sources of non-Gaussian error PDFs, chapter 4 considered a specific, geophysically relevant source, namely the mis-positioning of coherent features within a fluid domain. Chapter 4 argued for why coherent features should pose a problem to traditional methods of NWP and DA: in essence, NWP models and their DA systems take an Eulerian viewpoint to fluid evolution, whereas the position errors accrued by coherent features are more aptly treated in a Lagrangian viewpoint; *Gaussian Lagrangian errors can readily appear as non-Gaussian errors in an Eulerian representation*. Chapter 4 then considers some Lagrangian DA ideas to establish that standard ensemble-based filters can succeed in the face of coherent features when Lagrangian position information is explicitly taken into account. Such ideas are interesting and illuminating, but the methods considered are not likely to be useful to the NWP community where so much infrastructure has been built around Eulerian discretized forecast models. Hence, chapter 5 considered how one might begin to use the ideas from chapter 4 in an Eulerian context. The main idea explored there was the notion and utility of alternative error models, particularly error models containing alignment (i.e., position) errors. We showed that assuming such an alternative error model and adhering to state-of-the-art ensemble-based filter schemes naturally suggests a two-step approach to DA when coherent features are present. This was demonstrated in a simple proof-of-concept series of examples using a single, highly idealized, one-dimensional feature, the Korteweg-de Vries (KdV) soliton. The results are encouraging, but it remains to be seen if such an approach can be sensibly generalized to more realistic fluid flows where one may encounter multiple features and have to work in multiple spatial dimensions. This chapter addresses the generalization of chapter 5's results to more realistic scenarios.

This chapter is arranged as follows. We begin by presenting the methodology of augmenting state vectors, and discuss its power in an ensemble setting to bridge the elusive gap between Eulerian and Lagrangian (and perhaps even other) state information. We show how adherence to the state-of-the-art ensemble-based filters requires use of an alignment scheme to ensure that the analyzed augmented state vectors are internally consistent. This helps us re-interpret and generalize the two-step approach (and error model assumptions) from chapter 5. Another nice aspect of the augmented state vector approach is that it allows use of all available observations in each step as opposed to separating them as we did in chapter 5. After describing the augmented state vector approach, we then discuss the philosophy and mechanics of basis estimation. We focus on the design and mechanics of an alignment scheme fit for two-dimensional fluids. The alignment scheme we present

here is in no way purported to be an optimal, or even a great, method for accomplishing this, but the method does exhibit the main features that we argue an alignment scheme should have. Ultimately, our work is more about demonstrating the need for an alignment scheme rather than the specifics of a good scheme. Having decided on an acceptable alignment scheme, we present the results from two different observation system simulation experiments (OSSEs) using the proposed two-step augmented state vector approach with an intervening alignment scheme: *a*) interacting barotropic vortices (the dynamic analog of the point vortex experiments discussed in chapter 4), and *b*) a barotropically unstable jet breaking up into vortex patches. The OSSEs demonstrate that the proposed approach is much more effective than the state-of-the-art methods. The method’s success shows the power of ensemble-based approaches, even those restricted to being nearly linear and nearly Gaussian.

## 6.1 An augmented state vector approach

The Lagrangian estimation methods explored in chapter 4 indicate that explicitly embracing Lagrangian information can improve DA when coherent features are involved. In the case of point vortices, we found it is more effective to transform to a Lagrangian framework for DA purposes, that is, to estimate the vortices’ new positions and then move them there rather than allow the DA machinery to move the vortices by way of the discretized vorticity values. However, we also found that such an approach cannot generalize to larger, more realistic state vector estimates: one is not always able to cleanly transform an Eulerian problem into a Lagrangian one with significantly fewer degrees of freedom. A useful extension to this idea that can be applied to a more general class of state vectors is that of using augmented state vectors. This is a classic technique used throughout the control and state estimation literature to simplify a system (e.g., augmenting time to a state vector of a non-autonomous system to render the system “time independent” again) or to explicitly link quantities that have an otherwise complicated relation to state mathematically (e.g., simultaneous state and parameter estimation (Togneri and Deng 2003)). Here we use the technique to surmount the complicated relationship that exists between Eulerian field values and their corresponding Lagrangian position information.

The essence of this approach is to augment a state vector,  $\mathbf{x}$ , with the positions of all important features within its domain. As discussed in sections 4.2.4 and 5.3.2, the function that elicits position information from a state vector of a discretized field is nonlinear and cannot be sensibly linearized. One possible method of trying to find the position of a field’s

maximum value is to attempt to fit a very high order polynomial through all the discretized gridpoint values (essentially Lagrange interpolation), differentiate that polynomial, and then solve the resulting expression for its zeros (being careful to check to ensure a maximum has been reached rather than a minimum). This procedure may in principle be linearized, but we suspect it is a nearly impossible task for even moderately sized domain. Another, more implementable approach is to evaluate a search function like *arg max* that sifts through all values in a state vector and reports the index of the element with the highest value. Though easy to implement, this approach is not linearizable because of the discrete searching (its Jacobian entries are either 0 or  $\pm\infty$ ). Without a linearizable method to evaluate position, DA methods that rely on linearized operators (like the EKF and the adjoint method, see sections 2.3.3 and 2.4.2) will not be able to fully benefit from the augmentation, that is, they will not be able to have proper state dependent uncertainty estimates nor will they be able to use any observations bearing position information. This, however, is not an issue in an ensemble context.

### 6.1.1 Augmenting mechanics and statistics

Here we denote retrieved position information as  $\mathbf{p}$ , and it is found from the function  $\mathcal{F}[\cdot]$ . We then define an augmented vector  $\boldsymbol{\xi}$  such that

$$\boldsymbol{\xi} = \begin{bmatrix} \mathbf{x} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{x} \\ \mathcal{F}[\mathbf{x}] \end{bmatrix}. \quad (6.1)$$

This augmented state does not contain any new information *per se* since the augmented part is a function of the original state vector. However, its value becomes clear when one examines the covariance matrix corresponding to  $\boldsymbol{\xi}$ , particularly in an ensemble context.

Technically, the covariance matrix, here denoted by  $\boldsymbol{\Omega}$ , comes from the expected value of errors associated with  $\boldsymbol{\xi}$ , which is strictly determined only by the errors in  $\mathbf{x}$ . If  $\boldsymbol{\varepsilon}$  is a random error associated with  $\mathbf{x}$ , then the associated expected error in  $\mathbf{p}$  is related through a Taylor series of  $\mathcal{F}$  about  $\mathbf{x}$ :

$$\mathcal{F}[\mathbf{x} + \boldsymbol{\varepsilon}] \approx \mathcal{F}[\mathbf{x}] + \boldsymbol{\varepsilon} \frac{\partial \mathcal{F}}{\partial \mathbf{x}} + \dots, \quad (6.2)$$

but even if  $\boldsymbol{\varepsilon}$  is small in some appropriate manner, the Jacobian matrix of  $\mathcal{F}$  is not well-defined. Hence, it is not clear *a priori* how to specify  $\boldsymbol{\Omega}$  based on the expected errors in  $\mathbf{x}$  alone. Relatedly, it is also unclear how to specify the covariance between  $\mathbf{x}$  and  $\mathbf{p}$ . However, one of the advantages of an ensemble is that  $\boldsymbol{\Omega}$  need not be specified, it can simply

be approximated directly from an appropriately dispersed ensemble. In the ensemble case, each ensemble member can be diagnosed for its own  $\mathbf{p}$  using the full nonlinear  $\mathcal{F}$ , and then its variance and covariance can be directly estimated from the ensemble. We use the same symbol to denote the augmented ensemble, only subscripted:  $\boldsymbol{\xi}_j$ , for  $j = 1, N$ , where  $N$  is the number of ensemble members.  $\boldsymbol{\xi}_j$  is comprised of the state vector ensemble, written as  $\mathbf{z}_j$  as it was in earlier chapters, and the diagnosed position ensemble, written as  $\mathbf{p}_j$ . This means that the covariance matrix is (approximately):

$$\boldsymbol{\Omega} = \frac{1}{N-1} \sum_{j=1}^N (\boldsymbol{\xi}_j - \bar{\boldsymbol{\xi}}) (\boldsymbol{\xi}_j - \bar{\boldsymbol{\xi}})^T = \left[ \begin{array}{c|c} \text{cov}(\mathbf{x}, \mathbf{x}) & \text{cov}(\mathbf{x}, \mathbf{p}) \\ \hline \text{cov}(\mathbf{p}, \mathbf{x}) & \text{cov}(\mathbf{p}, \mathbf{p}) \end{array} \right], \quad (6.3)$$

where  $\bar{\boldsymbol{\xi}}$  is the ensemble mean value of  $\boldsymbol{\xi}_j$ . The matrix representation on the right shows how  $\boldsymbol{\Omega}$  is segmented. The ensemble not only allows us to find a satisfactory estimate of  $\text{cov}(\mathbf{p}, \mathbf{p})$ , but also, more importantly, we can find an estimate of  $\text{cov}(\mathbf{x}, \mathbf{p})$  (or equivalently, its transpose). When used within an estimation scheme, these cross-covariances potentially allow for innovations in one quantity to correct the other — this means that observing  $\mathbf{x}$  can contribute to correcting  $\mathbf{p}$ .

For clarity and completeness, we note that state augmentation provides the same uncertainty information to an ensemble-based Kalman filter as does using the “observed ensemble” technique described in sections 2.3.3 and 5.3.2: both are performing linear regression between  $\mathbf{x}$  and  $\mathbf{p}$ . The difference is that the augmented state vector approach folds the nonlinearity of  $\mathcal{F}$  into the state, whereas the other retains the nonlinearity within the observation operator  $\mathcal{H}$ . *One consequence of this difference is that the augmented state vector approach effectively increases the state size, and as we shall see, this is quite key since position becomes an explicitly treated quantity.*

## 6.1.2 Augmented state vectors and state estimation

### The “naive” approach

Proceeding forward, applying an ensemble-based Kalman filter to an ensemble augmented in the fashion described above, we might expect the cross-covariance information to constrain the estimation process and improve our analysis ensemble. One might even hope that by diagnosing the position information we have side-stepped the problems coherent features pose to estimation problems. *However, it is clear that this is not the case in general, for if  $\mathbf{z}_j^f$*

is initially non-Gaussian, then it still does not satisfy the basic assumptions of the ensemble filter. As stated in the previous section, the cross-covariances allow for the components to communicate in the correction process, but each component is still subject to the basic assumptions of the filter. A clearer explanation of this effect comes from considering a generic update of a prior ensemble  $\mathbf{z}_j^f$ . To form  $\boldsymbol{\xi}_j^f$ , one first evaluates  $\mathbf{p}_j^f = \mathcal{F}[\mathbf{z}_j^f]$ , where we have loosely extended our notation to indicate that  $\mathcal{F}$  is applied to each ensemble member. Hence,  $\boldsymbol{\xi}_j^f = [\mathbf{z}_j^{fT} \mid \mathbf{p}_j^{fT}]^T$ . Once some observations,  $\mathbf{y}^o$ , become available (unspecified for this generic case), we can perform DA:

$$\boldsymbol{\xi}_j^a = \boldsymbol{\xi}_j^f + \mathbf{K}(\mathbf{y}^o - \mathcal{H}[\boldsymbol{\xi}_j^f]) \quad (6.4)$$

$$\begin{bmatrix} \mathbf{z}_j^a \\ \mathbf{p}_j^a \end{bmatrix} = \begin{bmatrix} \mathbf{z}_j^f \\ \mathbf{p}_j^f \end{bmatrix} + \mathbf{K} \left( \mathbf{y}^o - \mathcal{H} \begin{bmatrix} \mathbf{z}_j^f \\ \mathbf{p}_j^f \end{bmatrix} \right). \quad (6.5)$$

DA has given a full analysis augmented vector,  $\boldsymbol{\xi}_j^a$ , comprised of analysis values of both components. As the Kalman filter is limited to corrections by linear combinations, it will not be able to satisfy the nonlinear relationship held between  $\mathbf{z}_j^a$  and  $\mathbf{p}_j^a$ . That is,

$$\mathbf{p}_j^a \neq \mathcal{F}[\mathbf{z}_j^a], \quad (6.6)$$

even though  $\mathbf{p}_j^f \equiv \mathcal{F}[\mathbf{z}_j^f]$ . The analysis is not internally consistent. *Said more explicitly, the analyzed feature positions,  $\mathbf{p}_j^a$ , need not, and in general will not, coincide with the feature locations in the analyzed state,  $\mathcal{F}[\mathbf{z}_j^a]$ .* This of course makes sense —  $\mathbf{p}_j^a$  implies that a feature is in a new position, and we know from section 5.3.3 and all the results of chapters 4 and 5 that the Kalman filter machinery is not able to reposition coherent features except when there are small position errors present. We have already seen an example of this in our cross-validation analysis in section 4.2.4; there the updated Eulerian state vector's analysis positions (the left panel of figure 4-8) hardly differed from its prior positions (the left panel of figure 4-5), whereas the Lagrangian update of those positions gives a large analysis increment (the left panel of figure 4-7).

### A two-step augmented state vector approach

This raises the question as to whether the augmented state vector analysis is useful at all beyond the traditional ensemble-based formulations. Simply put, yes, if the prior estimate satisfies the prior assumptions. If  $\mathbf{p}_j^f$  is nearly Gaussian, then  $\mathbf{p}_j^a$  should be a sensible, meaningful analysis, even though  $\mathbf{z}_j^a$  may not be. Actually, by this same reasoning, the

tails of  $\mathbf{z}_j^a$ , far away from the areas of position discrepancy should be a sensible, meaningful analysis as well, however, it is difficult to be selective about which updated state elements within  $\mathbf{z}_j^a$  one will accept, and even if one could make such a decision, it is unclear what to do with the elements where the analysis is deemed unacceptable. But as  $\mathbf{p}_j^a$  is a separate component altogether, accepting its analysis while rejecting  $\mathbf{z}_j^a$  is a cleaner exercise. To further clarify the development of the two-step method presented in chapter 5, we define two matrices that allow access to the different components of  $\boldsymbol{\xi}_j$ :

$$\mathbf{p}_j = \mathbf{S}_1 \boldsymbol{\xi}_j \quad (6.7)$$

$$\mathbf{z}_j = \mathbf{S}_2 \boldsymbol{\xi}_j, \quad (6.8)$$

where the subscripted numbers anticipate their matrices' uses in the two-step method. Each column of  $\mathbf{z}_j$  (i.e., each ensemble member) has  $n$  elements, and each column of  $\mathbf{p}_j$  has  $n_p$  elements. Hence,  $\mathbf{S}_1$  is an  $n_p \times n + n_p$  matrix comprised of the concatenation of an  $n_p \times n$  matrix of zeros and an identity matrix of dimension  $n_p$  (i.e.,  $\mathbf{S}_1 = [\mathbf{0}|\mathbf{I}]$ ).  $\mathbf{S}_2$ , an  $n \times n + n_p$  matrix, is similarly defined and constructed to access the  $\mathbf{z}_j$  component (i.e.,  $\mathbf{S}_2 = [\mathbf{I}|\mathbf{0}]$ ). Hence, the core mechanics of the two-step method can be written as:

$$\mathbf{p}_j^a = \mathbf{S}_1 \left( \boldsymbol{\xi}_j^f + \mathbf{K}_1 \left( \mathbf{y}^o - \mathcal{H} \left[ \boldsymbol{\xi}_j^f \right] \right) \right) \quad (6.9)$$

$$\mathbf{z}_j^a = \mathbf{S}_2 \left( \widehat{\boldsymbol{\xi}}_j^f + \mathbf{K}_2 \left( \mathbf{y}^o - \mathcal{H} \left[ \widehat{\boldsymbol{\xi}}_j^f \right] \right) \right), \quad (6.10)$$

where  $\mathbf{K}_1$  in equation (6.9) is the gain matrix calculated using  $\boldsymbol{\Omega}^f$ , the covariance matrix found using  $\boldsymbol{\xi}_j^f$  according to equation (6.3). Equation (6.10) requires some further comments. Because the Kalman filter machinery cannot reposition  $\mathbf{z}_j^f$  properly so that its features are in the positions specified by  $\mathbf{p}_j^a$ , we have had to suppose the intermediate state  $\widehat{\boldsymbol{\xi}}_j^f$  can be constructed, where  $\widehat{(\cdot)}$  indicates its argument has been *aligned*:

$$\widehat{\boldsymbol{\xi}}_j^f = \begin{bmatrix} \widehat{\mathbf{z}}_j^f \\ \mathbf{p}_j^a \end{bmatrix} \quad (6.11)$$

$$\mathbf{p}_j^a = \mathcal{F} \left[ \widehat{\mathbf{z}}_j^f \right]. \quad (6.12)$$

$\widehat{\mathbf{z}}_j^f$  still has a superscript  $f$  because  $\mathbf{z}_j^f$ 's values should not have been changed in constructing  $\widehat{\boldsymbol{\xi}}_j^f$ .  $\mathbf{K}_2$  is the gain matrix calculated using  $\widehat{\boldsymbol{\Omega}}^f$ , the covariance matrix found from  $\widehat{\boldsymbol{\xi}}_j^f$ .

The two-step philosophy is made very clear by equations (6.9) – (6.12). The error model considerations from chapter 5 are still very much present, they have just been absorbed into

the linear/Gaussian framework of the specific estimation scheme in which we have chosen to work. The augmented state vector approach makes the linear regression involved transparent. Working with  $\boldsymbol{\xi}_j^f$  allows one to easily approximate  $\text{cov}(\mathbf{x}, \mathbf{p})$ , while the operators  $\mathbf{S}_1$  and  $\mathbf{S}_2$  control which component one is updating. *Because we know the physical source of the non-Gaussianity in these situations, we are able to address it. Even though the PDF approximated by  $\mathbf{S}_2\boldsymbol{\xi}_j^f$  is non-Gaussian, the PDF approximated by  $\mathbf{S}_1\boldsymbol{\xi}_j^f$  can still be Gaussian. We have effectively found a transformation in this class of problems that renders the problems more Gaussian.* However, unlike traditional transformations to render problems more Gaussian (e.g., working with log-normal distributions for non-negative quantities), this transformation implies a two-step procedure: perform DA on  $\mathbf{S}_1\boldsymbol{\xi}_j^f$  so that  $\mathbf{S}_2\widehat{\boldsymbol{\xi}}_j^f$  is more Gaussian and can undergo successful DA.

This is a clear formulation of the problem, clear enough to highlight the questions and problems one will have to answer and solve to be able to implement this procedure in general. First and foremost, it remains to be elucidated how one constructs  $\widehat{\mathbf{z}}_j^f$  from knowledge of  $\mathbf{z}_j^f$  and  $\mathbf{p}_j^a$ . This is a potentially large subject, and we will address the philosophy and mechanics of this estimate realignment in the next section. Also, we will address the existence and construction of the operator  $\mathcal{F}[\cdot]$ .

## 6.2 Basis estimation and alignment schemes

This section considers the philosophy and mechanics of changing a state estimate so that it adheres to newly specified position information, that is, finding  $\widehat{\mathbf{z}}_j^f$  from  $\mathbf{z}_j^f$  and  $\mathbf{p}_j^a$  (symbols defined in section 6.1.2). We refer to the DA problem that generates  $\mathbf{p}_j^a$  as *basis estimation* since in the parlance of alignment error models it is suggesting how one should modify the gridpoint locations,  $s_i$ . We refer to the process that generates  $\widehat{\mathbf{z}}_j^f$  as an *alignment scheme* since it enforces that the forecast ensemble reflect the two-step method's efforts to estimate alignment errors. *An alignment scheme is necessary to complement the two-step method because, while ensemble-based Kalman filters can successfully produce analysis feature positions,  $\mathbf{p}_j^a$ , they cannot produce corresponding state estimates whose analyzed features are positioned at  $\mathbf{p}_j^a$ .* Hence, in order to capitalize on the reduction in alignment errors from DA, one needs to enforce that the features within  $\mathbf{z}_j^f$  become positioned at  $\mathbf{p}_j^a$ . We note that this was assumed in chapter 5's presentation of the two-step method where the alignment scheme was a trivial global shift (assuming periodic boundaries) of a one-dimensional profile such that  $p^f$  was mapped to  $p^a$  (see figure 5-6). If one is to extend this two-step method to states in multiple spatial dimensions and/or bearing multiple features,

then one must implement a non-trivial alignment scheme. We offer one such alignment scheme based on an image processing technique called Piecewise Linear Mapping, but we stress that this is not necessarily the best scheme one could implement. We begin this section by discussing the notion of basis estimation, and we distinguish between the basis used in model space (physical space) and in state space. We then discuss qualities that the ideal alignment scheme would exhibit, followed by a presentation and discussion of our method based on piecewise linear mapping.

### 6.2.1 Basis estimation

One interpretation of the alignment error model written in equation (5.4) is that there are errors in  $s_i$ , the locations of the gridpoints. As we mentioned in section 5.3.3, it is perhaps a bit counterintuitive to allow for errors in the gridpoint locations since they are often regarded almost as a computational axiom, but here we discuss why this interpretation arises and why it is an acceptable interpretation.

#### The philosophy of basis estimation

Within state estimation, there are two different spaces of interest, and hence there are two different bases of concern:

- 1.) State space is spanned by the *vector basis*, which serves to logically order the elements with respect to one another (e.g., indexed neighbors like  $\mathbf{x} = [\dots x_{i-1} x_i x_{i+1} \dots]^T$ ).
- 2.) Discretized physical space, or model space, is spanned by the *spatial basis*, which affixes labels to model gridpoints to place and orient the state vector elements in physical space.

A given state vector is a point in its state space, and a discretized, spatially-extended field in its model space. The vector basis always has dimension  $n$ , whereas the spatial basis will have as many elements as there are gridpoints and require as many labels per element as there are spatial dimensions. For example, if  $\mathbf{u} = [\mathbf{v}^T \mathbf{w}^T]^T$  (a bivariate state vector) and each of  $\mathbf{v}$  and  $\mathbf{w}$  are represented over a 2D grid with coordinates at  $(x, y)_i$ , where  $i = 1, n$ , then the spatial basis has dimension  $n$  and requires two labels, whereas the vector basis has dimension  $2n$ , one for each element in  $\mathbf{u}$ .<sup>1</sup> Being clear on these differences is important as

---

<sup>1</sup>Note that specifying the gridpoints as  $(x, y)_i$  for  $i = 1, n$  implies that each gridpoint has its own two labels. This is in the spirit of Matlab's "meshgrid" command or in the spirit of having irregular grid spacing. If the  $x_i$ 's and  $y_i$ 's are regularly spaced, then one could resort to a Cartesian coordinate system of  $(x_j, y_k)$ , where  $j = 1, n_x$ ;  $k = 1, n_y$ ; and  $n_x \times n_y = n$ .

vector calculus operations like gradients are applied over the spatial basis, whereas linear algebra operations like finding the mean and covariance are applied over the vector basis.

If the evolution of a state vector,  $\mathbf{x}$ , includes coherent feature positioning errors (whose distinction from other errors only makes sense in model space), then one may find it convenient to question the “axiom” of having precisely known gridpoint locations. Of course, the gridpoints’ locations are known and fixed by virtue of the forecast model’s numerics, but since estimation schemes like the Kalman filter cannot in general correct position errors on their own, directly correcting position errors by allowing for uncertain gridpoint locations is much more approachable than trying to get the same effect from somehow consistently correcting all additive errors at the corresponding gridpoint locations. *Stated another way, in certain cases, estimation of the spatial basis itself may be more sensible than that of the vector elements directly.* Note that as the spatial basis is related to the vector basis by way of labels (coordinates with units of distance), one can imagine adjusting the label values in an effort to better align the model representation of the fluid field with that thought to be truth’s. And while one can alter the spatial basis, the vector basis remains untouched in the sense that the state vector will always have  $n$  elements ordered along monotonically increasing spatial coordinate axes.

The first step of the two-step approach is to perform spatial basis estimation. To justify adjusting the gridpoint labels at all, one needs a prior estimate of the labels, that estimate’s uncertainty, and some observations that give information about the “true values” of the labels. From the point of view of the forecast model’s numerics, all estimates share the same exact locations of gridpoints, and the prior estimate of their locations and its uncertainty are  $\mathbf{s}_i$  itself with zero uncertainty. *However, one way to capitalize on the usefulness of spatial basis estimation and the two-step method is to begin viewing feature position error as errors in the gridpoint locations.* For instance, within an ensemble of fluid state forecasts, if every member has a differently positioned representation of a feature, then invoking the belief that each member’s feature is an estimate of the same true feature (i.e., drawing correspondence between features) gives an ensemble estimate of the uncertainty in that feature’s position. The feature’s position is thus a “tie-point,” an identifiable marker within the field that helps characterize the estimate’s morphological landscape, that is, the field’s shape (e.g., the positions of highs and lows or ridges and troughs). The ensemble allows approximation of a tie-point’s uncertainty, thus allowing DA to better estimate the tie-point’s position. The ensemble analysis positions of the tie-point allows one to “remap” the spatial basis so that a given ensemble member’s current feature position is relabeled in space as the analysis position.

Once the analysis tie-point locations are known, then the specifics of how to relabel all other gridpoints are left to a particular alignment scheme (this will be covered in the next section). There are other ways to approach aligning an estimate, but we stress that it is only through tie-points that we can access information about spatial basis *uncertainty* to close a well-posed DA problem. Fixed boundaries constitute another class of tie-points, but these tie-points truly have zero uncertainty since boundaries are set locations in space (i.e., all ensemble members agree on the boundaries' locations). Therefore, one only has cause to begin relabeling the spatial basis if one can identify interior tie-points, and this inevitably means identifying features (we will briefly address the feasibility of this issue in the next section). Since we have already acknowledged that the gridpoint locations are, in fact, known exactly by virtue of the model numerics, it is clear from the outset that any adjustment done to the spatial basis labels by DA will have to be compensated for by some spatial interpolation of the aligned estimate back onto the original gridpoint locations before the estimate can be integrated further by the model.

To summarize, if one has an ensemble of states,  $\mathbf{z}_j^f$ , from which one can diagnose the positions of a feature,  $\mathbf{p}_j^f$ , then one can approximate the uncertainty in  $\mathbf{p}_j^f$  from the ensemble spread. If the diagnosed features all correspond to the same true feature, then  $\mathbf{p}_j^f$  contains  $N$  estimates of a tie-point. If observations are available that contain any information on the true location of the tie-point, then  $\mathbf{p}_j^f$ , its uncertainty, the observations, and their uncertainty close a DA problem which can be handled by an ensemble-based Kalman filter.<sup>2</sup> Solving the DA problem produces the analysis positions,  $\mathbf{p}_j^a$ , but DA cannot move the features in  $\mathbf{z}_j^f$  so that they are located at  $\mathbf{p}_j^a$ . Consider ensemble member  $j = 1$ : one way to move  $\mathbf{p}_1^f$  to  $\mathbf{p}_1^a$  is to simply relabel the gridpoint at  $\mathbf{p}_1^f$  with the value specified by  $\mathbf{p}_1^a$ . If there are no other tie-points, then one has no other information about the spatial basis's uncertainty, and one must assume all gridpoints shift globally in accordance with the  $\Delta\mathbf{s}$  implied by  $\mathbf{p}_1^a - \mathbf{p}_1^f$ . To have only one tie-point implies that the boundaries must be periodic, as fixed boundaries are by definition tie-points. In the case of multiple tie-points, it is clear the remapped basis spacing will in general change throughout the domain as a single global shift will not remap all tie-points to their analysis positions. The details of how the spacing changes are left to a specific alignment scheme, and we shall discuss several implementations in the next subsection.

---

<sup>2</sup>Of course, the same caveats apply to the involved PDFs being nearly Gaussian.

## Schematic tie-point driven alignment example

A schematic example of remapping a spatial basis guided by tie-point analysis is shown in figure 6-1. The top panel shows a hypothetical one-dimensional forecast state estimate to be regarded as a particular ensemble member. It is comprised of a Gaussian “hump” with an included background sine wave to highlight the possible effect of an alignment scheme bringing an estimate “out of balance.” The state estimate is demarcated by regularly spaced black circles to show how the spatial basis resolution can change as a result of the alignment scheme. The red diamond denotes a hypothetical analysis position to which the first step of DA has determined this ensemble member’s feature should be remapped. The middle panel shows how the alignment scheme would align the state estimate when considering only one tie-point, that is, with periodic boundaries. Without any other information to guide the spatial basis remapping, the alignment scheme moves all gridpoints uniformly so that the forecast position is now at the analysis position. This is precisely the manner in which the alignment scheme worked in the soliton examples in chapter 5 (see figure 5-6). The bottom panel shows a possible aligned state estimate from considering three tie-points, the feature and two fixed boundaries (at  $s_1 = -10$  and  $s_n = 10$ ). All alignment schemes would need to shrink resolution in the left half of the domain and stretch resolution in the right half in order to satisfy the three tie-point constraints, but alignment schemes are free to differ on how they choose to fill-in the gridpoint spacing between tie-points. This example simply linearly spaces the points between the tie-points, but many possibilities are imaginable.

Because one global shift is insufficient for satisfying all three tie-points, distortion has arisen, evident mainly in the background sine wave. One can imagine how such distortion might upset balance in a physically meaningful example; however, all DA methods short of a particle filter can potentially bring state estimates out of balance.<sup>3</sup> That the state estimate requires an alignment scheme is an indication that its errors have a non-Gaussian PDF, and we have already demonstrated that not using an alignment scheme can cause significant distortion in its own right. The important point is that the features are being moved on the basis of a trusted DA problem.

### 6.2.2 Alignment schemes

In this subsection, we briefly consider the attributes that an ideal alignment scheme would have (if one existed) and a few possible approaches for approximating this ideal scheme.

---

<sup>3</sup>If one considers smoothers, then a so-called “strong constraint” method will give an analysis that is guaranteed to be a model solution.

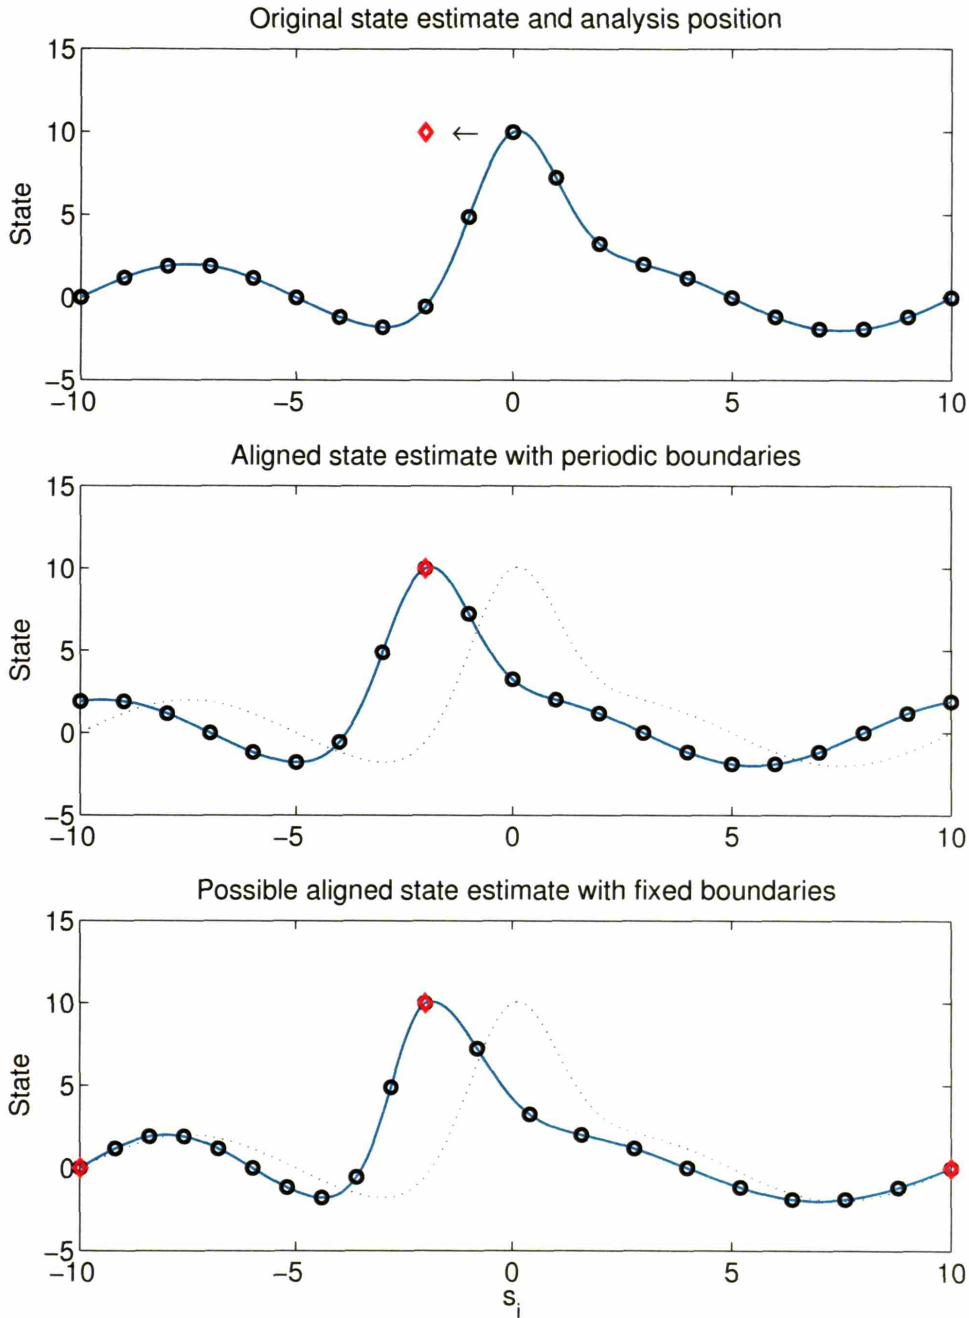


Figure 6-1: The top panel shows a hypothetical forecast state estimate and an analysis position (red diamond) that we want our estimate's feature to match. The black circles are evenly-spaced markers to show the spatial basis distortion. The middle panel shows what an alignment scheme would produce based on a single tie-point. The bottom panel shows a possible realigned state from considering three tie-points, the feature's position and the two fixed boundaries.

## Ideal alignment

The ideal alignment scheme would reposition all involved tie-points with the minimal amount of distortion to the rest of the field, where the meaning of “distortion” needs to be defined and quantified. One attempt to define distortion in this setting is the subject of a classic problem in what is today called operations research, namely the problem of optimal mass transport, commonly called the Monge-Kantorovich problem (e.g., Evans 1997). Monge, a forefather of differential geometry, originally posed the problem in the 1780s. He was interested in the way to move a given mound of dirt to fill a specified hole that requires the minimum amount of work. In the 1940s, Kantorovich, an economist and the inventor of linear programming, revisited the problem while studying the optimal allocation of resources. The problem is now generally posed as searching for the diffeomorphic, measure-preserving mapping from one distribution to another that minimizes some specified cost-function (Evans 1997). A diffeomorphism is a specific kind of mapping generally referred to as “one-to-one and onto,” meaning that all points in one distribution uniquely map into the second and vice versa, implying that diffeomorphisms are invertible. A measure over a distribution or part of a distribution is a non-negative real number that allows one to meaningfully compare distributions’ sizes. Probability is a commonly encountered example of a measure. For a mapping to be measure-preserving means that the mapping does not change a distribution’s size (e.g., probabilities still integrate to unity). When a diffeomorphic, measure-preserving mapping exists between two distributions, it is rarely unique, and the Monge-Kantorovich problem seeks the *optimal* mapping such that a cost-function is minimized. Monge originally considered a cost function that penalized the absolute value of distance ( $\ell^1$ -norm) for each mapped bit of the distributions, and Kantorovich considered a cost function that penalized the square of these distances ( $\ell^2$ -norm) (Evans 1997). Either minimizing mapping could constitute a mapping of minimal distortion if distortion is linked to the cost function each is minimizing. Haker et al. (2004) go further to pose a slightly different definition of what they term the minimal distortion mapping, though it is closely linked to their proposed solution method of Monge-Kantorovich problems.

The Monge-Kantorovich problem has been studied in many fields and has been used in widely varying applications, most notably in the problem of image registration and warping (Zhu et al. 2003; Haker et al. 2004). To quote Haker et al. (2004), “image registration is the problem of establishing a common geometric reference frame between two or more image data sets possibly taken at different times” (page 225). This sounds very similar to the stated goal of an alignment scheme, and it is particularly pleasing as it is based on solutions to the Monge-Kantorovich problem which boasts minimum distortion in some

defined sense. However, the applicability of the work of Haker et al. (2004) or others is not immediately clear because our approach is different. The Monge-Kantorovich problem solves for a mapping based on two completely specified distributions without any references to tie-points or other landmark data. This does not capture our problem since we are trying to use the alignment scheme (or Monge-Kantorovich solution) to find the second distribution. This means we are seeking the solution to a different problem that is tie-point driven. Working with tie-points means that our registration problem has already taken place — in order to approximate the uncertainty of the tie-points, we had to assert that they referred to the same true feature. It is unclear to us whether one could derive an alignment scheme based on merging the Monge-Kantorovich problem with estimation methods, and we did not attempt to carry this proposition too far as it is beyond the scope of our work. However, we hold out hope that one may be able to devise an alignment scheme that supplies a minimal work, diffeomorphic, measure-preserving mapping of a forecast ensemble (where some measure must first be defined, as a fluid field is not itself a measure) to an analysis state that satisfies supplied tie-point information. We note that the Field Alignment algorithm of Ravela et al. (2004), which implements regularization by local constraints of smoothness and non-divergence, seems a good first order approximation of the Monge-Kantorovich solution, and is a step in an encouraging direction.

In the next subsection, we present a possible alignment scheme. We are aware that this matter is beyond our expertise, and we merely try to offer a plausible scheme that accomplishes the specified alignment without causing egregious distortion of the fields being aligned. We note that the details of the alignment scheme may not be crucial since the second step of our approach can potentially ameliorate offending distortions.

### **Alignment scheme example**

A two-dimensional alignment scheme is conceptually just a generalization of the one-dimensional examples shown in figure 6-1. Because two spatial dimensions requires two labels per gridpoint, we introduce further notation to more easily distinguish vector bases from spatial bases. We retain our established notation for the vector basis, that is,  $i$  is an index over state elements (whereas  $j$  is an index over ensemble members). Since  $i$  is an index over the vector basis, it references elements stored in a column vector, not on the discretized grid. Therefore, the index  $i$  has little regard for the actual two-dimensional geometry of the spatial basis (i.e., the spatial relation between elements  $i_o$  and  $i_o + 1$  is not immediately obvious). Discussion of alignment schemes is generally clearer from the standpoint of the spatial basis, since that is the basis being updated. In general, the geometry of the gridpoints

in two-dimensions can be completely arbitrary, however, in most reasonable fluid prediction scenarios, we expect a somewhat regular or at least predictable geometry. If the gridpoints are arranged in a rectangular grid, as many models choose to do, then the  $n$  gridpoints can be referenced in a two-dimensional grid fashion where there are  $n_x \times n_y (= n)$  gridpoints. We define two new indices:  $k = 1, n_x$  and  $l = 1, n_y$ . Hence,

$$\mathbf{s}_i = \mathbf{s}_{k,l} = (s_x, s_y)_{k,l}, \quad (6.13)$$

and now the spatial relation between elements  $(k_o, l_o)$  and  $(k_o + 1, l_o + 1)$  is immediately clear. For a univariate state vector (where  $n$  is also the number of gridpoints), the vector basis is  $\mathbf{s}_i$ , and the spatial basis is  $\mathbf{s}_{k,l}$ . This same distinction between bases applies to diagnosed tie-point positions as well:  $\mathbf{p} = (p_x, p_y)$ .

To be more concrete about the distinction between these two bases, we introduce the following example: consider the alignment of an interior tie-point within a two-dimensional domain with fixed boundaries. We affix simple numbers to this example so that the reader does not get confused by the notation, and we refer to this example throughout this subsection. We consider an initial interior tie-point position,  $\mathbf{p}^f = (p_x^f, p_y^f) = (4, 4)$ , and an analysis position,  $\mathbf{p}^a = (p_x^a, p_y^a) = (6, 5)$ . These tie-point locations happen to be a gridpoint locations, but they need not be. Since an alignment scheme adjusts the spatial basis according to  $\mathbf{p}^a$ , we will distinguish the spatial basis before and after application of the alignment scheme by the same superscripts,  $f$  and  $a$ . In this example, we consider a spatial basis with relatively coarse resolution, for the purpose of demonstration. We set  $n_x = n_y = 9$  and assume the domain is square with initially uniform grid spacing of 1 in both coordinate directions. This means that there is a good deal of redundancy in the values of  $\mathbf{s}_i^f$ , but since the alignment scheme can alter the gridpoints' labels, there is not expected to be much redundancy in  $\mathbf{s}_i^a$ . Explicitly, one arrangement of  $(s_x, s_y)_i$  into a column vector is:

$$\begin{aligned} s_{x,i} &= [1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7 \ 8 \ 9 \ 1 \ \dots \ 9]^T \\ s_{y,i} &= [1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 2 \ \dots \ 9]^T. \end{aligned}$$

This arrangement implies a rule for converting between the vector basis index and the spatial basis index:  $i = n_y(l - 1) + k$ . Hence,  $\mathbf{s}_1 = \mathbf{s}_{1,1} = (1, 1)$  and  $\mathbf{s}_{10} = \mathbf{s}_{1,2} = (1, 2)$ . By this convention, for our chosen points,  $\mathbf{p}^f = \mathbf{s}_{4,4}^f = (4, 4)$  and  $\mathbf{p}^a = \mathbf{s}_{6,5}^a = (6, 5)$ . All alignment schemes will produce an updated spatial basis such that  $\mathbf{s}_{4,4}^a = (6, 5) = \mathbf{s}_{6,5}^f$  and all the fixed boundary points remain unaltered. Alignment schemes are free to differ in how each finds the elements of  $\mathbf{s}^a$  that are not tie-points, however, they should all abide by the

relative ordering of elements in the vector basis and the arrangement rule for transforming to the spatial basis. This means that we should not allow an alignment scheme to switch elements as would be the case if, say,  $\mathbf{s}_{3,3}^a = \mathbf{s}_{7,7}^f$  and  $\mathbf{s}_{7,7}^a = \mathbf{s}_{3,3}^f$ .

One possible aligned basis for this example is shown in figure 6-2. The left panel shows the initial spatial basis, and the right panel shows the aligned spatial basis. In the left panel,  $\mathbf{p}^f$  is shown by the blue diamond and  $\mathbf{p}^a$  is shown by the red star — the aim is to shift the gridpoint shown by the blue diamond to the location marked by the red star and have all the other gridpoints move in a sensible way. As is evident in the right panel, the employed alignment scheme has simply performed linear interpolation in both coordinate directions. This method is chosen just to highlight the main feature of the alignment, namely that the gridpoint corresponding to  $\mathbf{p}^f$  has been mapped so that its position is at  $\mathbf{p}^a$  and that the surrounding gridpoints with no direct tie-point information must be remapped accordingly. We next describe the alignment that we developed for the proof-of-concept examples considered later in this chapter.

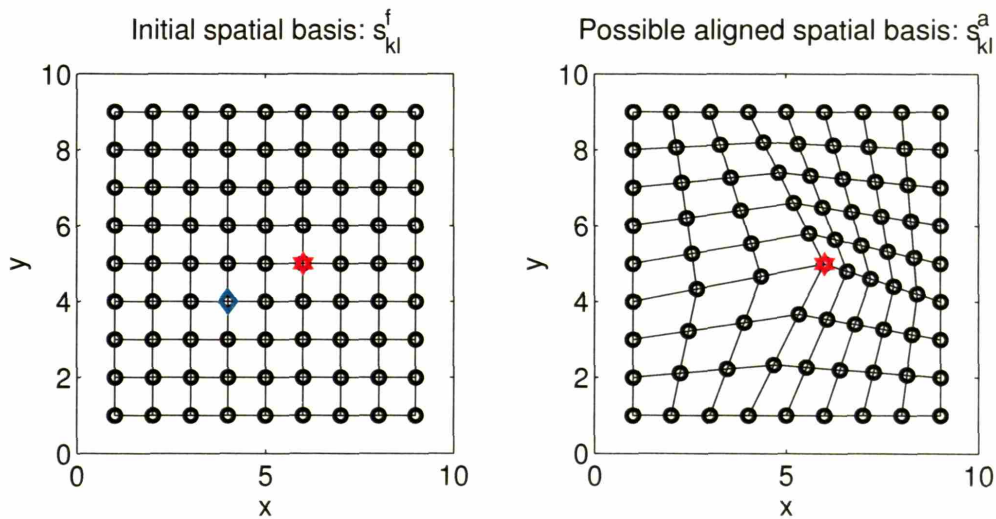


Figure 6-2: The left panel shows the initial spatial basis. The black circles are the gridpoints, the blue diamond shows  $\mathbf{p}^f$ , and the red star shows  $\mathbf{p}^a$ . The right panel shows a possible realigned basis. The black lines show the same neighbor relations that the gridpoints have in the left panel.

### Piecewise linear mapping

While an alignment scheme based on the Monge-Kantorovich problem would be ideal, it is unclear whether one can be constructed. Without demanding the ideal scheme, we (merely)

seek an alignment scheme that is general enough to handle an arbitrary number of tie-points, is free of user-defined parameters, and tends to spread grid distortion over a “necessarily affected” region (e.g., somewhere between the entire grid and just the local neighborhood). One major difference between a hypothetically ideal alignment scheme based on the Monge-Kantorovich problem and the one described here is that the former would perform alignment based on the values of the fluid field (e.g., in an effort to be measure-preserving), whereas ours just considers mechanical transformations on the spatial basis itself.

We have selected a technique based on the piecewise linear mapping technique first described by Goshtasby (1986), where one uses a special geometric relationship between the tie-points in the original spatial basis to segment the basis into smaller regions that are then mapped individually into the updated basis guided by the analysis positions of the tie-points. The special geometric relationship is called the *Delaunay triangulation*. The Delaunay triangulation is the dual of the *Voronoi tessellation* (sometimes also called the Dirichlet tessellation), and both arise in optimally characterizing a discrete set of points in a plane.<sup>4</sup> The edges delineated by the Delaunay triangulation define so-called natural neighbor relations for point sets (Watson 1992).

The Delaunay triangulation and its corresponding Voronoi tessellation have many special properties, and can thus be derived in many ways. Perhaps the most intuitive is to consider the Voronoi tessellation: it is the result of dividing a plane into polygonal “tiles,” where each tile contains a single tie-point within its area. The tiles’ boundaries are such that all other points within the tile are closer to the tie-point within that tile than any other tie-points in the plane. Hence, the Voronoi tessellation naturally divides the plane into “regions of influence” for each tie-point. The Delaunay triangulation results from connecting tie-points whose tiles share a common edge. The Voronoi tile edges thus bisect the the Delaunay triangulation edges. To make this connection clearer, the Voronoi tessellation and Delaunay triangulation for the tie-points of the  $9 \times 9$  spatial basis example are shown in figure 6-3. The left panel shows the Voronoi tessellation. There are 33 tiles shown, one for each tie-point. The blue lines show the tile edges. The outer bounding square exists here because the domain is bounded, but technically when considering the infinite plane, the lines normal to the bounding square should extend to infinity. The right panel repeats the Voronoi tessellation in the thin blue lines and the Delaunay triangulation is overlain in red. Note that all tie-points with contiguous tiles are connected. This means that the corner points are not connected to the interior tie-point. Note also that the blue tile edges of the interior

---

<sup>4</sup>The concepts immediately generalize to three dimensions as well where one has the Delaunay “tetrahedralization.”

tie-point bisect the red lines connecting the interior point to the boundaries. The Delaunay triangulation defines each tie-point's natural neighbors.

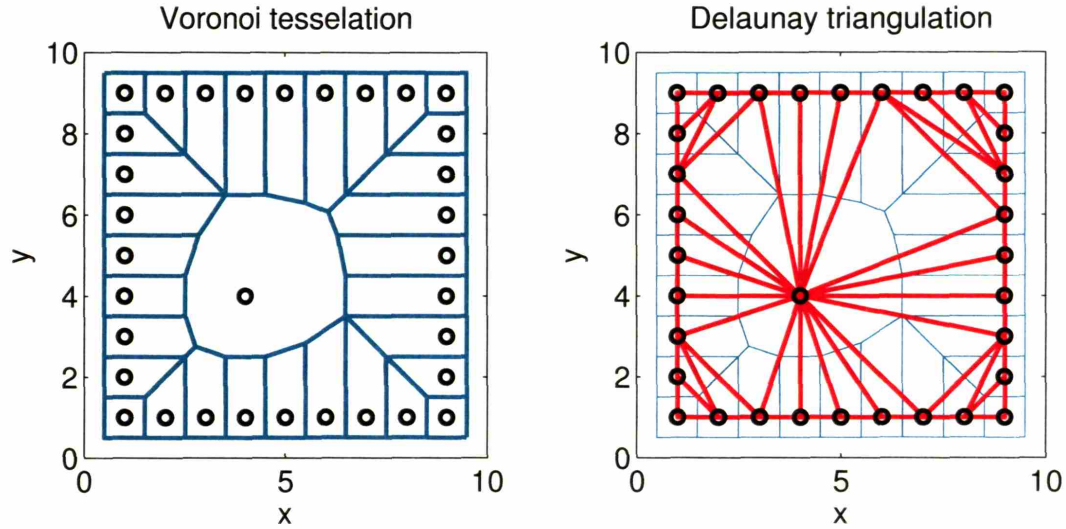


Figure 6-3: The left panel shows the Voronoi tessellation for the tie-points shown by the black circles. The right panel additionally shows the Delaunay triangulation in red.

Once the Delaunay triangulation of the tie-points has been found in the initial spatial basis,  $\mathbf{s}_i^f$ , then the natural neighbor relations have been established. Piecewise linear mapping assumes that the tie-points have the same natural neighbor relations in the aligned basis,  $\mathbf{s}_i^a$ , as in the initial spatial basis. This means the tie-points have the same triangulation before and after alignment, though the triangulation may not necessarily be the Delaunay triangulation of  $\mathbf{s}_i^a$ . Retaining the same triangulation allows one to segment both spatial bases into triangular tiles (these are different than the Voronoi tiles). *The crux of piecewise linear mapping is to transform all gridpoints in the initial spatial basis into the aligned basis based on their triangular tile membership.* To transform the gridpoints in one triangle to another, one can simply perform an *affine transformation*. Affine transformations preserve collinearity and ratios of distances, though not necessarily angles or lengths (e.g., Weisstein 1999). Subsumed transformations in affine transformations include translation, scaling, rotation, and shearing. The general affine transformation of one point's location,  $(p_x^f, p_y^f)$ , to another,  $(p_x^a, p_y^a)$ , is:

$$p_x^a = A + Bp_x^f + Cp_y^f \quad (6.14)$$

$$p_y^a = D + Ep_x^f + Fp_y^f, \quad (6.15)$$

where  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $E$ , and  $F$  are constants. The affine transformation is thus determined by three points (6 equations for the 6 unknown constants). This means that all triangles are affine, that is, any triangle can be transformed into any other triangle. Hence, the vertices of the triangular tiles are sufficient for determining how all interior gridpoints should be transformed.

The piecewise linear mapping for our example is shown in figure 6-4. The left panel shows the same Delaunay triangulation as in the right panel of figure 6-3, but here plotted with all gridpoints, not just the tie-points. Also, the analysis interior tie-point position is shown by the red star. The right panel shows the same triangulation for the analysis tie-point positions. Note that this triangulation is not the Delaunay triangulation of the analysis tie-points. The initial interior tie-point position is shown by the blue diamond. Note how the corresponding triangles still contain the same number of gridpoints, and how the affine transformations of each triangular tile have retained the gridpoints' initial collinearity and ratios of distances.

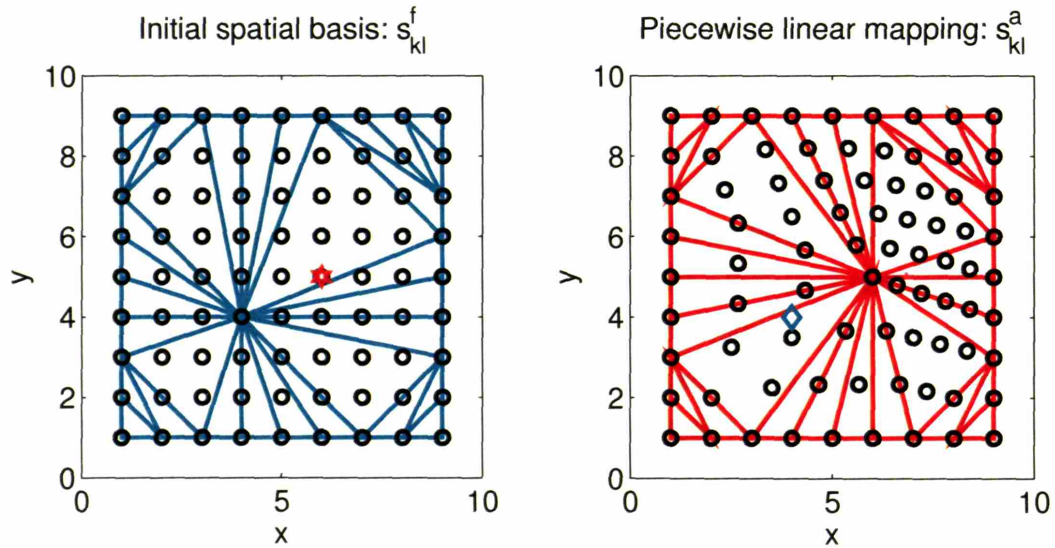


Figure 6-4: The left panel shows the initial spatial basis,  $s_{kl}^f$ , and its Delaunay triangulation based on its tie-points. The red star denotes the analysis position of the interior tie-point. The left panel shows the aligned spatial basis from applying piecewise linear mapping. The blue diamond shows the initial position of the interior tie-point.

For comparison with the very simple scheme shown in figure 6-2, figure 6-5 shows an overview of the piecewise linear mapping alignment scheme's results. The top left panels shows the initial spatial basis with gridlines drawn. As before, the blue diamond shows the initial interior tie-point's position, and the red star shows its analysis position. The upper

right shows the same aligned spatial basis as shown in right panel panel of figure 6-4, only here the gridlines are drawn instead of the triangulation edges. Though piecewise linear mapping does not operate on each coordinate direction separately, the bottom two plots are included to show how the natural neighbor relationships spread information. Note the alignment scheme’s hesitancy to develop strong angles in the gridlines, especially away from the interior tie-point. Note also that the corner points whose triangulations do not connect them to the interior tie-point are not altered at all. This indicates that the Delaunay triangulation tends to move only those gridpoints “that must be,” where the natural neighbor relationships define this concept.

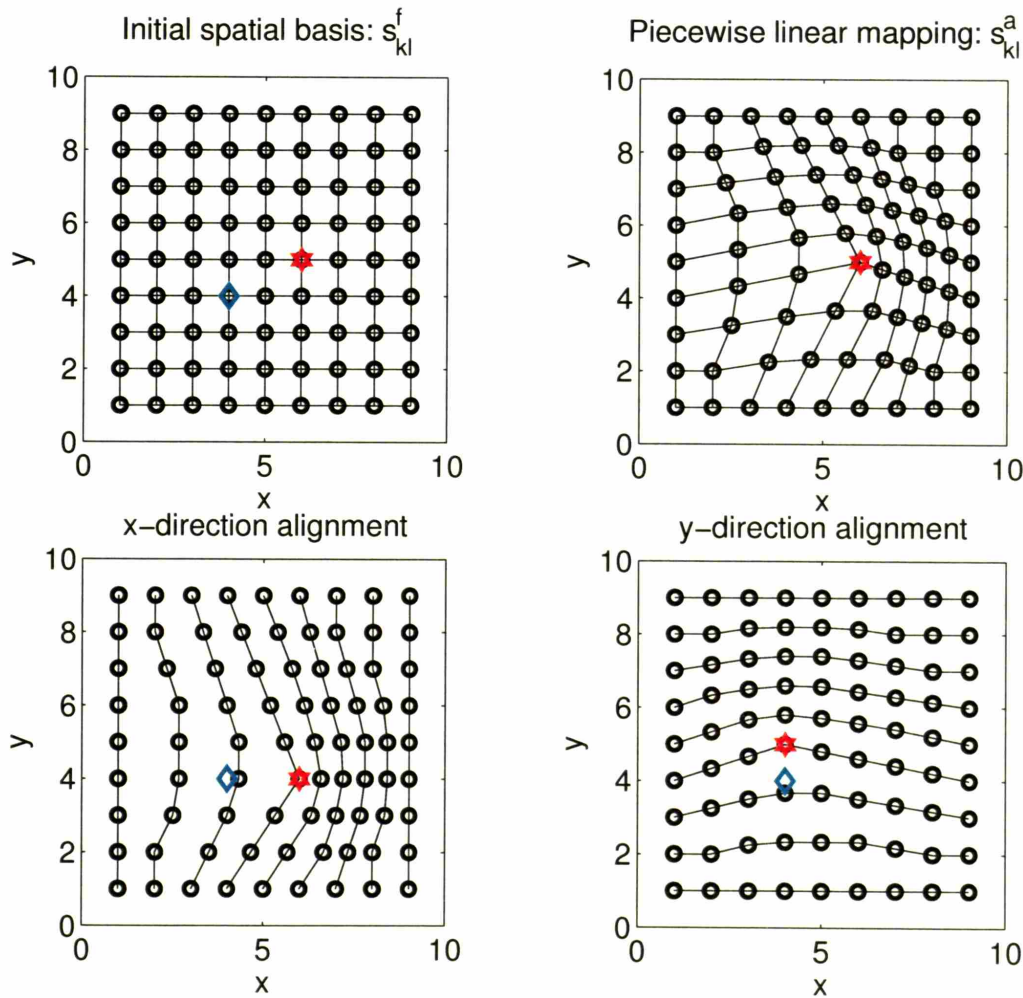


Figure 6-5: The top two panels show the same information as figure 6-2: the top left shows the initial spatial basis with the blue diamond showing  $\mathbf{p}^f$  and the red star showing  $\mathbf{p}^a$ ; the top right shows the aligned spatial basis via piecewise linear mapping; the bottom left shows the  $x$ -coordinate update and the bottom right shows the  $y$ -coordinate update.

It should be noted that piecewise *cubic* mapping exists as well (Goshtasby 1987). The advantage of piecewise cubic mapping over the linear mapping is analogous to using cubic splines over linear interpolation, that is, the obtained mapping is smoothly differentiable as well. We did not choose to implement the more involved cubic algorithm as the added complexity does not grossly change the resulting alignment scheme, and it therefore seems inconsistent with our notion of just finding an *adequate* scheme (i.e., we are aware that better schemes surely exist, better than the cubic scheme as well). We also note that the Delaunay triangulation has already found its way into modeling fluid evolution via the discipline of unstructured grids (e.g., Legrand et al. 2000).

### 6.2.3 Feature identification & tracking

In this subsection, we briefly address the operator  $\mathcal{F}[\cdot]$  from section 6.1 that is responsible for eliciting tie-point information from gridded fluid state estimates. Our general treatment of  $\mathcal{F}$  is simply to assume it exists. By our hesitancy to define a feature, there is likely to be some subjectivity in identifying one. However, there are several points of which to remain aware in this matter. The first is that no matter how exotic  $\mathcal{F}$  is in form, it can certainly be evaluated because the user has the entire specification of each ensemble member available. In principle,  $\mathcal{F}$  diagnoses a state estimate,  $\mathbf{x}$ , to find  $\mathbf{p}$ , and insofar as any  $\mathbf{p}$ 's can be found, the information in  $\mathbf{x}$  should be sufficient for their diagnosis. We add that feature recognition is a very large field with its own canon of literatures — if a sophisticated  $\mathcal{F}$  is required for implementation of a (pseudo) operational form of the two-step method, then further research will certainly be required into this foreign literature. The  $\mathcal{F}$  operators we use for the OSSEs shown in this thesis are generally the simple *arg max* operators referred to in sections 4.2.4 and 5.3.2, though in the more complicated experiments we have sometimes been forced to use manual feature identification. Also, as reviewed in section 4.1.2, sophisticated methods of feature identification and tracking have already made their way into the atmospheric sciences, motivated by the field of model validation for feature-based climate signals like storm tracks. Hodges in particular has been a proponent of applying pattern recognition techniques to general circulation model (GCM) output and satellite images (Hodges 1994, 1995, 1998, 1999; Hoskins and Hodges 2002; Anderson et al. 2003; Delsol and Hodges 2003; Delsol 2005). He and his coauthors have offered automated procedures for both feature identification procedures and feature tracking procedures.

Another matter which helps the implementation of  $\mathcal{F}$  is the fact that the features of interest are presumed coherent. This means that in practice one will more often be updating one's estimate of  $\mathbf{p}$  than diagnosing it for the first time. This coherency through time

suggests that so-called feature trackers can be of use. Feature trackers are widely used throughout science for intermittently observed states. Aside from the Hodges references above, examples in the earth sciences include supercell storm trackers used over relatively short time scales (i.e., nowcasting) (Wilson 2004), tracking water vapor imagery features to derive so-called “satellite winds” (Velden et al. 1997), and tracking ice sheet features with periodic sweeps of synthetic aperture radar (Luckman et al. 2003).

### 6.2.4 Feature correspondence across the ensemble

In regards to the correspondence problem, that is, what to do if not every ensemble member has a feature or the same number of features, we note that an ensemble of large enough membership should be able to handle this without too much trouble. The main crux is that the ensemble of diagnosed positions  $\mathbf{p}_j$  need not have the same size as the ensemble of states,  $\mathbf{z}_j$ , whence they came. The first step of the two-step procedure need only include those members whose features correspond to one another. One should not become too worried by the absence of a feature in some ensemble members as its absence contributes to the uncertainty of the state, and will be reflected in the second step of the two-step procedure.

## 6.3 OSSEs

With an acceptable alignment scheme in hand, we are in a position to apply the two-step methodology to more realistic problems than the proof-of-concept KdV solitons considered in the previous chapter. More specifically, we would like to demonstrate that this approach can be applied to a fluid in two spatial dimensions, possibly bearing multiple features. In this section we consider two such scenarios, both using the same model, a barotropic model. In the subsections that follow, we briefly describe the model, and then we present results from two different model set-ups: interacting point vortices and a jet going unstable. For each of the model set-ups, we run observation system simulation experiments (OSSEs) similar to those shown in chapters 3 and 5. In both cases, the OSSEs compare the standard ensemble-based DA approach with our two-step approach.

### 6.3.1 A barotropic model

A barotropic fluid is one whose dynamics can be encapsulated by a single variable, either the vorticity,  $q$ , or equivalently, a streamfunction,  $\psi$ . This allows treatment of a single

evolution equation, the *barotropic vorticity equation* (BVE). On a rotating plane, the BVE can be written:

$$\frac{\partial q}{\partial t} = -J(\psi, q) - \beta \frac{\partial \psi}{\partial x} + \mathcal{D} + \mathcal{N} + \mathcal{R}, \quad (6.16)$$

where  $J$  is the Jacobian operator,  $q = \nabla^2 \psi$ ,  $\beta$  is the ambient meridional vorticity gradient,  $\mathcal{D}$  is an unspecified form of dissipation,  $\mathcal{N}$  is an unspecified form of numerical dissipation (small scale damping), and  $\mathcal{R}$  is an unspecified form of forcing, often written as a relaxation term.

In the model runs that follow, we assume  $\beta$  is always zero (i.e.,  $f$ -plane dynamics only),  $\mathcal{D}$  is always zero (i.e., no physical sources of friction like Rayleigh damping), and that  $\mathcal{R}$  is always zero (i.e., no forcing), leaving the ensuing dynamics to be driven by self-consistent vorticity dynamics and numerical dissipation. The term  $\mathcal{N}$  is necessary for numerical stability. We implement two different versions of  $\mathcal{N}$ , one for each of the cases treated, the details of which are given in appendix D. The particular forms used will be mentioned in each of the following subsections. The given spatial and time resolutions will also be addressed in each subsection. As we have interest in maintaining accurate representation of coherent features, we choose to implement a so-called “pseudospectral” method to solve equation (6.16) with doubly periodic boundary conditions, a third order Adams-Bashforth time-stepping routine (Durran 1991), and we implement standard two-thirds Nyquist wavenumber de-aliasing. The barotropic model used in this section is different than the one used in section 3.4.2: *we constructed our own model for these experiments because accurately representing coherent features in a simplified model such as a barotropic model requires some special care.*

A barotropic fluid is in many ways idealized. We employ its use here mainly because it is rich enough to support some resemblance of fluid features and yet computationally efficient enough to run economically many times for experimentation purposes, particularly given our ensemble approach. We note that there is a long tradition of using barotropic models for predictability studies, beginning with Charney et al.’s original attempt at NWP (1950). An Euler fluid, as treated in section 4.2.1, is an example of a barotropic fluid. More generally, a barotropic fluid can be considered an approximation of any system whose core dynamics admit a so-called “master variable” like potential vorticity. Though one cannot formally prove that synoptic-scale atmospheric dynamics occur on a separable “slow manifold” (Lorenz and Krishnamurthy 1987), there is wide agreement that the synoptic development with which forecasters are concerned is mainly contained within the atmosphere’s potential vorticity dynamics (e.g., Hoskins et al. 1985).

### 6.3.2 The point vortex model revisited

The first set of experiments we present concerns the representation of point vortices within a barotropic model. In section 4.2.4, we presented kinematic examples of updating point vortex state vectors and their corresponding Eulerian, gridpoint representations. Here we present a dynamic example, where the barotropic model is used to evolve the vorticity distributions between updates.

#### Inserting point vortices into a barotropic model

To properly evolve a point vortex population, a point vortex model should be used (see section 4.2.1). Inserting a  $\delta$ -function representation of  $q$  into a gridpoint based fluid model is of course problematic. To accomplish this feat, we resort to the Gaussian vortex approximation described in section 4.2.3. To remind the readers, the Gaussian vortex approximation (in cylindrical coordinates about a vortex center) is:

$$q(r) = A \exp(-Br^2) \tag{6.17}$$

$$v_\theta(r) = \frac{A}{2Br} (1 - \exp(-Br^2)), \tag{6.18}$$

where  $A$  and  $B$  are chosen so that  $v_\theta$  appears to be the azimuthal velocity of a point vortex outside of a certain range (i.e., if the Gaussian vortices are separated far apart enough, they will interact as point vortices). Specifying that  $q$  be one tenth its maximum value a distance  $d$  away from the center, we find  $B = -\frac{\log(0.1)}{d^2}$  and  $A = \frac{B\Gamma}{\pi}$ , where  $\Gamma$  is the circulation strength of the corresponding point vortex.

To successfully simulate point vortices with a barotropic model, one needs to judiciously balance the gridpoint resolution and the vortex scale  $d$  so that the vortices can be stably represented by the spectral numerics and still be far enough away from one another to interact at distance as point vortices. A lower bound on  $d$  is set by the necessary numerical dissipation, and this means that our model runs are sensitive to the form of  $\mathcal{N}$  we choose. For this case, we employ an exponential cut-off wavenumber filter, which leaves all wavenumbers above a certain length scale untouched and transitions to complete elimination of small scales; such a scheme was used by Arbic and Flierl (2003). See appendix D for details. Using this implementation, we find that to maintain coherent Gaussian vortices, we must use a  $d$  above  $2.5\Delta x$ , where  $\Delta x$  is the gridpoint resolution (assumed uniform in  $x$  and  $y$  directions over the whole grid). The minimum  $d$  we use is  $3\Delta x$ , giving the vortices an effective diameter of  $6\Delta x$ .

As discussed in section 4.2.1, one needs at least four point vortices to obtain a nonintegrable Hamiltonian (i.e., a system that exhibits sensitivity to initial conditions without having an attractor) for a point vortex model. Here with the doubly period boundaries, the existence of one vortex in the domain of interest implies an infinite number of vortices at regularly spaced distances away. However, because they are regularly spaced (by model domain lengths), they cannot interact chaotically (i.e., the degenerate cases referred to by Aref and Pomphrey (1982)). We find we still need to use at minimum four Gaussian vortices to obtain the “interesting” behavior considered in chapter 4, and this combined with the minimum Gaussian vortex core size determines the minimum number of gridpoints required to successfully simulate point vortices with a barotropic model. We find that in practice, a minimum of  $128 \times 128$  resolution is required as  $64 \times 64$  can still lead to close vortex interactions and vortex mergers.<sup>5</sup>

### Comparing to a doubly periodic point vortex model

In section 4.2.4, we compared the Lagrangian update of the point vortex state vector with its corresponding Eulerian state vector. If we are to make the same comparison in this dynamic example, then there is need to generalize the point vortex model to one that is doubly periodic. This is an important distinction because double periodicity implies a fixed domain size, where the original point vortex model has no explicit boundedness. It is also a matter of some computational complexity as there are infinite number of vortices to consider. We are not aware of any means to accomplish this analytically via a recursion relation or a covering limit. It is difficult to approximate this numerically by, say, explicitly adding a number of “shells” around the center domain that replicate it. The elliptical nature of the Laplacian operator makes each point vortex’s influence felt at a far distance, particularly when there is a large amount of cancelation between contributing terms. For example, in a doubly periodic square domain of length  $L$  centered at the origin, one ought to be able to position two vortices of equal strength at  $(\pm L/4, 0)$  and see no further change in their positions because of all the cancelations occurring from equally distant neighboring vortices. However, if one approximates the truly doubly periodic domain by, say,  $N_s$  shells of domains about the center one, there will be one unpaired vortex at either end whose effects cannot be canceled. Though the initial advecting velocities at the center vortices may be vanishingly small, the advecting velocities of the surrounding vortices will not be, and it will not take long for their effects to be felt in the center domain. This will set the vortices in the center domain into motion about one another. This will be true even for *very* large values of  $N_s$ ,

---

<sup>5</sup>The restriction to powers of 2 in the resolution is necessitated by the pseudospectral solution method.

though it is easiest to visualize with just one shell. Without presenting too much detail here, we have devised a highly accurate approximate solution to the doubly periodic point vortex model by transforming each vortex in the original domain into three “sub-vortices” whose strengths sum to the original vortex’s strength. See appendix E for more details.

Initializing both our doubly periodic point vortex model and our barotropic model with the same vortex positions, we are able to integrate them quite a long time (several error doubling times) before their positions noticeably diverge from one another.

### **Doubly periodic piecewise linear mapping**

Since the models we are filtering are doubly periodic, it is worth addressing how one might apply piecewise linear mapping to the states. Essentially, one must consider nearest neighbor relationships through boundaries. Doubly periodic models have no fixed boundary tie-points like those shown in the alignment scheme figures above in section 6.2.2. As mentioned in that section, if there is only one tie-point with doubly periodic boundaries, an alignment scheme should just circularly shift the entire estimate to accomplish the remapping. With more than one tie-point, one needs to consider the tie-points in the neighboring periodic domains for calculating the Delaunay triangulation. An example of this is shown in figure 6-6. The center domain is bounded by the heavy dashed line, and it contains four tie-points. The center domain’s periodic neighboring domains have been plotted as well, and each of those domains contain the same, regularly space tie-points. The superset of all nine domains’ tie-points then define a new Delaunay triangulation, shown by the blue lines. This can be considered a doubly periodic Delaunay triangulation, as the shaded gray triangle indicates. The analysis positions of all tie-points have been determined by the first step of the two-step method, so piecewise linear mapping can proceed normally. Afterwards, we only retain the gridpoints within the center domain’s area.

### **Experimental set-up**

As discussed in section 4.2.3, the point vortex model admits a non-dimensional parameter, therefore allowing identical results to be obtained at any scales. Instead of working with “geophysically relevant scales” (e.g.,  $O(10^3\text{km})$ ), we instead consider generic units of length and time (the results can always be scaled to whatever lengths are desired). Our goal is to run an OSSE with three parallel ensemble integrations: 1) a doubly periodic point vortex system filtered by a standard EnKF, 2) a similarly configured barotropic vortex system filtered by a standard EnKF, and 3) another identically initialized barotropic vortex system but filtered by our augmented state vector two-step approach with use of an intervening




Figure 6-6: This shows the doubly periodic Delaunay triangulation of the original four tie-points in the center domain, bounded by the thick dashed lines. The gray triangle shows the triangle connecting three tie-points through the period boundaries.

piecewise linear mapping alignment scheme. To accomplish this, there is need for the point vortex model and the barotropic model to be in time-synch, meaning we need to calibrate their time scales. Anticipating the barotropic model experiments run in the next subsection with the unstable jet, we choose to configure the point vortex model as follows: a square domain of length  $L = 25.6$  centered at the origin, five equal vortex strengths of  $\Gamma = L/4$ , and a time step of  $\Delta t_{pv} = 0.664$ . We have generally run the model using 4 shells, which is still accurate enough to keep correspondence over several assimilation cycles (see appendix E for more details).

To address the time scale synchronization, we used the doubly periodic point vortex model in this configuration with 64 shells to approximate its error doubling time. We find  $\tau_{\text{doub}} \approx 96\Delta t_{pv}$ . As described in section 4.2.3, this suggests the calibration of defining one “day” to be  $48\Delta t_{pv}$ . The complementary barotropic model configuration has the same

model domain size and initializes its Gaussian vortices with the same value of  $\Gamma$ , however, its numerics are more delicate and require a much shorter timestep. As will be discussed at length in the next subsection, Flierl et al. (1987) have used a similar barotropic model toward a similar feature-oriented end. We use their suggested value for  $\Delta t_{bt}$  of 0.0166 for use with a resolution of  $128 \times 128$ . This value is observed to abide by numerical stability conditions. Also, assuming the same error doubling time as the point vortex model (i.e., error growth is due to mis-positioning), we can calibrate the barotropic model’s time scale, finding that a “day” in the barotropic model is  $1920\Delta t_{bt}$ .

Knowing the error doubling time suggests a parameter range in which to run our OSSEs, however, it also depends crucially on the observation system. We have run the point vortex system in isolation many times to carve out the observational parameters where OSSEs are successfully able to filter the system (i.e., no filter divergence). As we are interested in distinguishing our two-step approach from the traditional ensemble-based Kalman filters, we choose to run our OSSEs in a range just prior to when the point vortex model fails, though we have confirmed that all three ensembles perform well as expected for very frequent, dense, and accurate observation systems.

The kinematic update results shown in section 4.2.4 were for a  $64 \times 64$  Eulerian model resolution, whereas here we are forced to use  $128 \times 128$  (meaning a state vector size of  $n = 16,384$ ). As in those experiments, we will assume that there are fixed station observations of fluid velocity available at a fixed time interval (called the observation time) as well as one position observation (though there are five vortices present). We always observe the same feature in these experiments as if we were tracking its development in particular. To avoid any sources of potential model error, each model has its own freely evolving version of truth that is observed, though as explained above, the point vortex system is able to track the barotropic model for quite some time. To avoid their eventual divergence we slightly adjust the point vortex system’s truth state at each assimilation to better agree with the barotropic model’s. This ensures that the various ensembles have enough information to track with each other, if their respective DA systems allow them.

There is a large parameter space to explore as well as a huge amount of information generated from each OSSE due the number of different ensembles being integrated. We make no attempt to fully describe and explore these models’ respective parameter ranges. *What we do wish to communicate is the following: a) there are of course observational parameter ranges where all the filtered ensembles perform correctly, b) there are observational parameter ranges where none of the ensembles are able to maintain reasonable estimates of truth, and c) the respective levels in parameter space where the ensembles begin to fail are*

*different.* On this last point, by the results in section 4.2.4, the point vortex system is of course able to filter truth most robustly since it represents the best case scenario for this experiment. Plus, it has only errors in position, whereas the barotropic ensembles need to retain accurate amplitudes for their vortices to maintain good representation of truth. *We find that the two-step method is able to extend the range of successful filtering over the EnKF's ensemble when each are given the same observational information taken from the same truth profile — for accurate enough observations, the two-step ensemble can approach equivalency with the point vortex ensemble because the first step of the two-step approach achieves most of the correction necessary.*

The following results are from an OSSE with an observation time of 18 “hours,” a station velocity observation error of 10% of the maximum velocity in truth’s profile (set by the Gaussian vortex profile), and a position observation error of one gridpoint. There are 50 randomly selected station velocity locations (confined to be gridpoints to alleviate the necessity of including interpolations within our observation operator). We use an ensemble size of  $N = 50$  for all three ensembles. For initial generation, the point vortex ensemble was generated first by the method described in section 4.2.3, and then the other two ensembles were found from the point vortex ensemble by evaluating the corresponding vorticity distributions under the Gaussian vortex assumption (i.e., precisely how it was done in section 4.2.4). We note that the point vortex DA begins to fail for these parameters at an observation time of 24 “hours,” so we are nearing the expected predictability limit for the information at hand. For the figures shown below, we have assumed a Gaussian vortex radius of  $5\Delta x$ . As indicated in section 4.2.4, the results are somewhat sensitive to this choice, and we comment on this after presenting the results here.

## OSSE results

The two-step ensemble is observed to track truth well for many assimilation cycles, whereas the EnKF ensemble loses correspondence to truth rather early on. We first present representative “overview” figures of the fourth assimilation cycle, where the EnKF’s individual members begin to lose physical plausibility. As discussed in section 3.3.2 and in chapter 5 in relation to the KdV experiments, an ensemble can still be accurate even when it has lost reliability, and we see this again for the EnKF’s ensemble mean after DA has been performed. After the overview figures, we trace one particular member through several successive assimilation cycles to better appreciate how the two-step method is able to maintain its accuracy while the EnKF fails. We conclude with a discussion and an “executive overview” figure comparing truth and the three ensemble means after the end of assimilation cycle.

### *Overview figures*

We begin by discussing the barotropic model's EnKF ensemble. Figure 6-7 shows the EnKF ensemble information prior to DA. The top left panel shows the barotropic truth state (the colored contours, though the contours are tightly packed enough so that they appear continuously colored), the 50 station velocity observations, and the black star over the vortex near the right edge denotes the position observation. As the model is doubly periodic, the vortices can go through the boundaries. So as to not obscure information at the edges, in all the ensuing figures we have demarcated the domain edges by black dashed lines. The upper right panel shows the EnKF forecast ensemble mean just prior to DA. The contour interval is the same in all panels (with the zero contour omitted and negative vorticity contours denoted by black dotted lines), so note the severe smearing and weakening that has occurred. The black dots denote the ensemble's diagnosed feature position of the observed feature. The dispersion of these dots gives a good indication of the spread of the ensemble members. However, to make it a bit clearer, we show two particular ensemble members in the bottom two panels. The lower left was subjectively deemed a "good" ensemble member for the EnKF ensemble, where good means that five well formed vortices agree with fairly closely truth's vortices' positions. The lower right was chosen as an example of a "bad" ensemble member. Note that the vortices have begun to get sheared away and have begun filamenting. The vortex whose position has been observed is still firmly intact, as is the vortex in the northwestern quadrant of the domain. From the ensemble mean (in the upper right panel) we see that this northwestern vortex is the strongest signal in the ensemble. Inspection of the "synoptic situation" reveals that it is a very slowly moving vortex at this point as most of the "action" is occurring with the three vortices on the eastern half of the domain encircling one another.

Figure 6-8 shows similar information to figure 6-7, only for the EnKF analysis states after DA has been performed. The upper left panel again shows truth and the available observations (the same as the upper left panel of figure 6-7). The upper right panel shows the EnKF analysis ensemble mean. It is clear the vortex signatures have been made sharper within the mean, indicating the strength of the observational content. However, inspection of the individual ensemble members shows that truth's vortex signal is only clearly imprinted within the mean: "good" and "bad" members alike have been updated somewhat "noisily." This the same sort of behavior we observed for the KdV ensemble states in section 5.5.3 (*cf.*, top panel of figure 5-11). One may rightfully note that the ensemble mean shown in the upper right panel of figure 6-8 is a good approximation of the minimum variance estimate

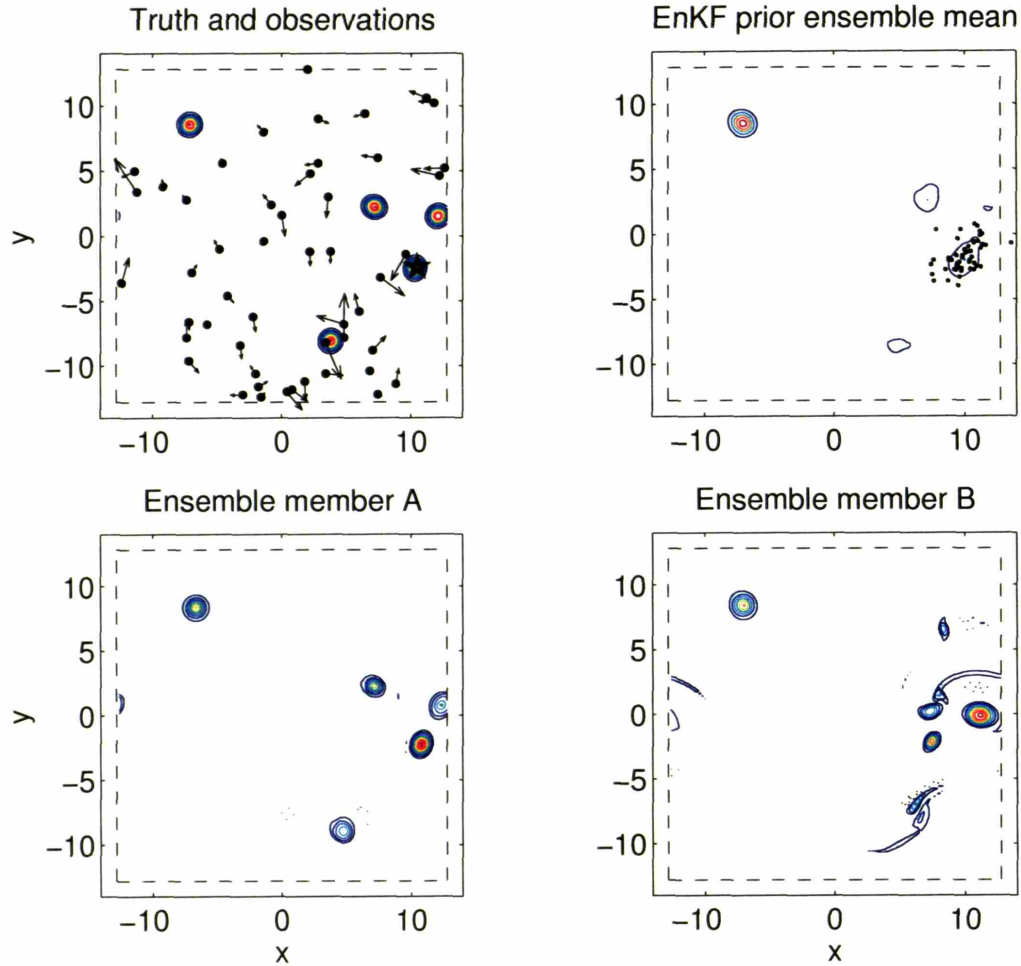


Figure 6-7: The top left shows the truth state of the barotropic model after four assimilation cycles. The colored contours show positive values of vorticity. The filled black circles denote station observation locations, and the black vectors show the velocity observations that have become available. The black star over one of the vortices denotes the position observation that has been made available. The black dashed border shows the model domain boundaries. The upper right panel shows EnKF ensemble mean prior to DA. All panels have the same contour interval, with the zero contour omitted and negative contours shown by thin black dotted lines. The black dots overlaying the mean are the diagnosed vortex positions from each ensemble member corresponding to the vortex whose position was observed. The lower two panels each show a different individual ensemble member.

of the state. However, its utility in ensemble-based filtering is dubious — the ensemble will not serve well as initial conditions for the next DA cycle, nor will it allow for instructive probabilistic interpretation.

We now examine the two-step method’s ensemble. In contrast to the ensemble shown

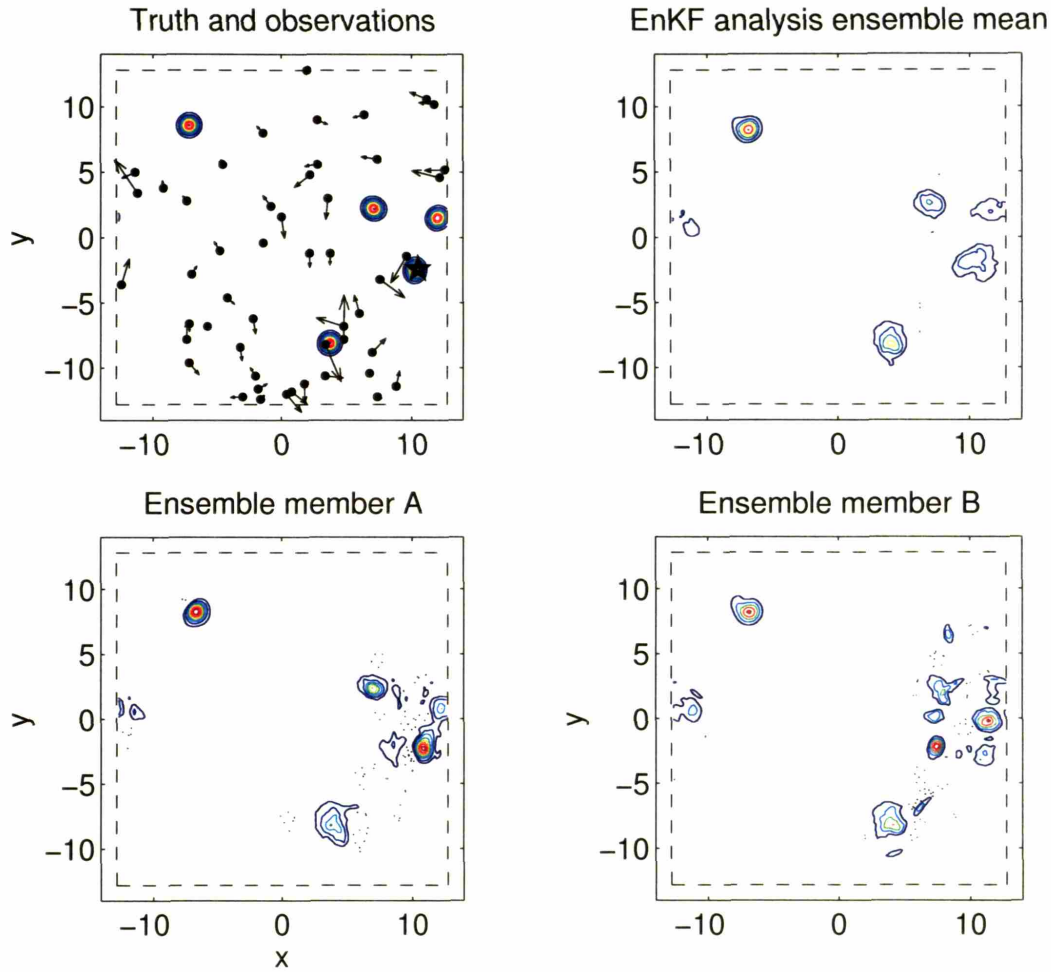


Figure 6-8: The four panels here show the same information as shown in figure 6-7, only for the EnKF analysis ensemble states. The contour interval and convention is also the same in this figure.

in figure 6-7, the two-step method's forecast ensemble, shown in figure 6-9, still shows good correspondence to truth, both in its mean and in the individual members. The information shown in figure 6-9 is analogous to the previous two figures. The upper left panel again shows truth and the observations. The upper right panel shows black dots for all five features now, because in the two-step approach we are first doing DA on all designated feature positions (though we are only observing one feature), whereas in the EnKF we only needed the positions of the observed vortex for computing the innovations. Note that this ensemble has received no additional information from truth (in the form of observations) over the EnKF ensemble: the diagnosed positions plotted constitutes information obtained from the ensemble. The two individual ensemble members plotted no longer constitute subjective

determinations of “good” and “bad,” they merely correspond to the same members plotted in figures 6-7 and 6-8 (they are worth comparing since they were equal when the initial ensembles were generated).

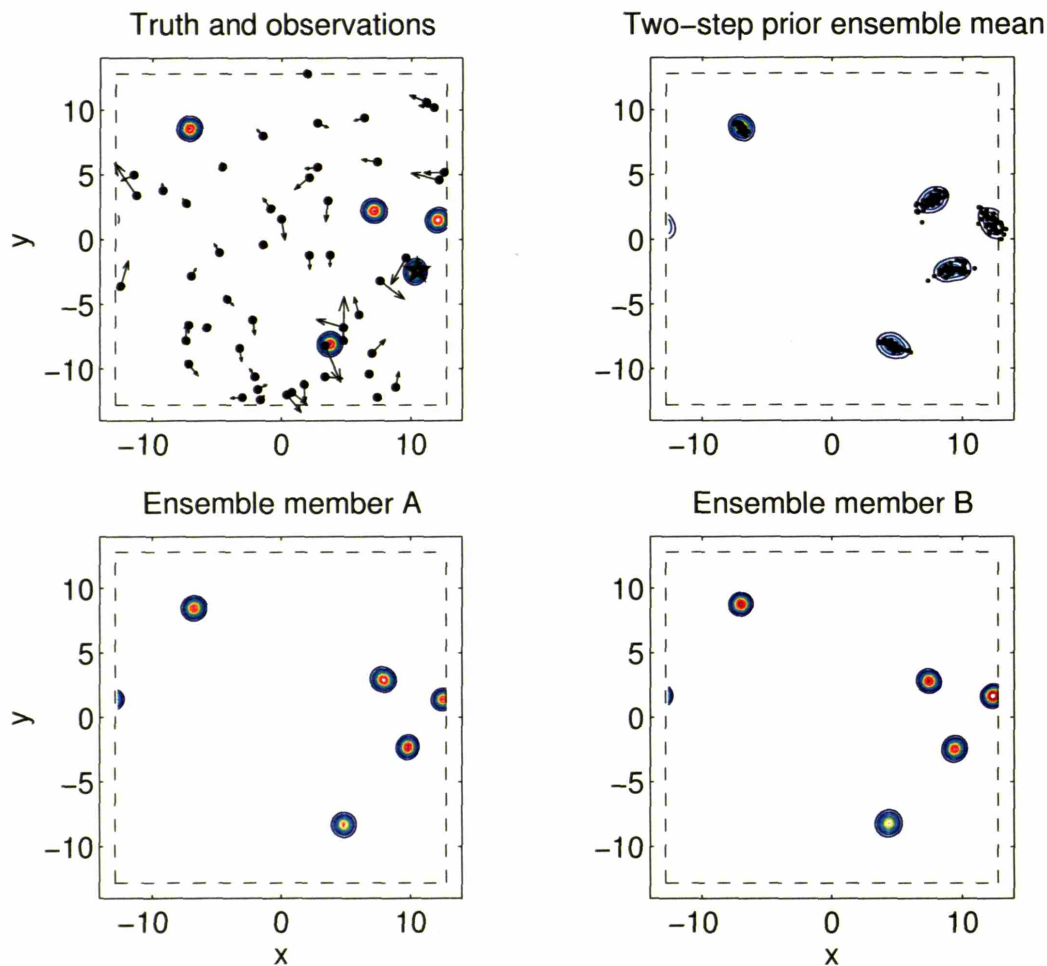


Figure 6-9: The four panels here show the same information as shown in figure 6-7, only for the two-step ensemble states prior to DA. Again, the contour interval and convention is the same as before.

The results of the first step of the two-step approach and of the alignment scheme are shown in figure 6-10. The upper left panel again shows truth and the observations. The upper right panel shows the ensemble mean after undergoing the alignment scheme, and the black dots overlain are the feature analysis positions from the first step of DA. Note that not all vortex positions had their uncertainty reduced equally — the two vortices with largest remaining spread are the most active vortices and the farthest from station observations. Even so, most barotropic model ensemble members retain five reasonably shaped vortices

in nearly the correct locations. The two individual members shown in the bottom panels confirm that the alignment scheme is not too “violent” to the members, at least when they need not be moved very far. Again, it is worth reiterating that the information used to produce this aligned ensemble was available to the EnKF ensemble as well, but the EnKF was not able to access it.

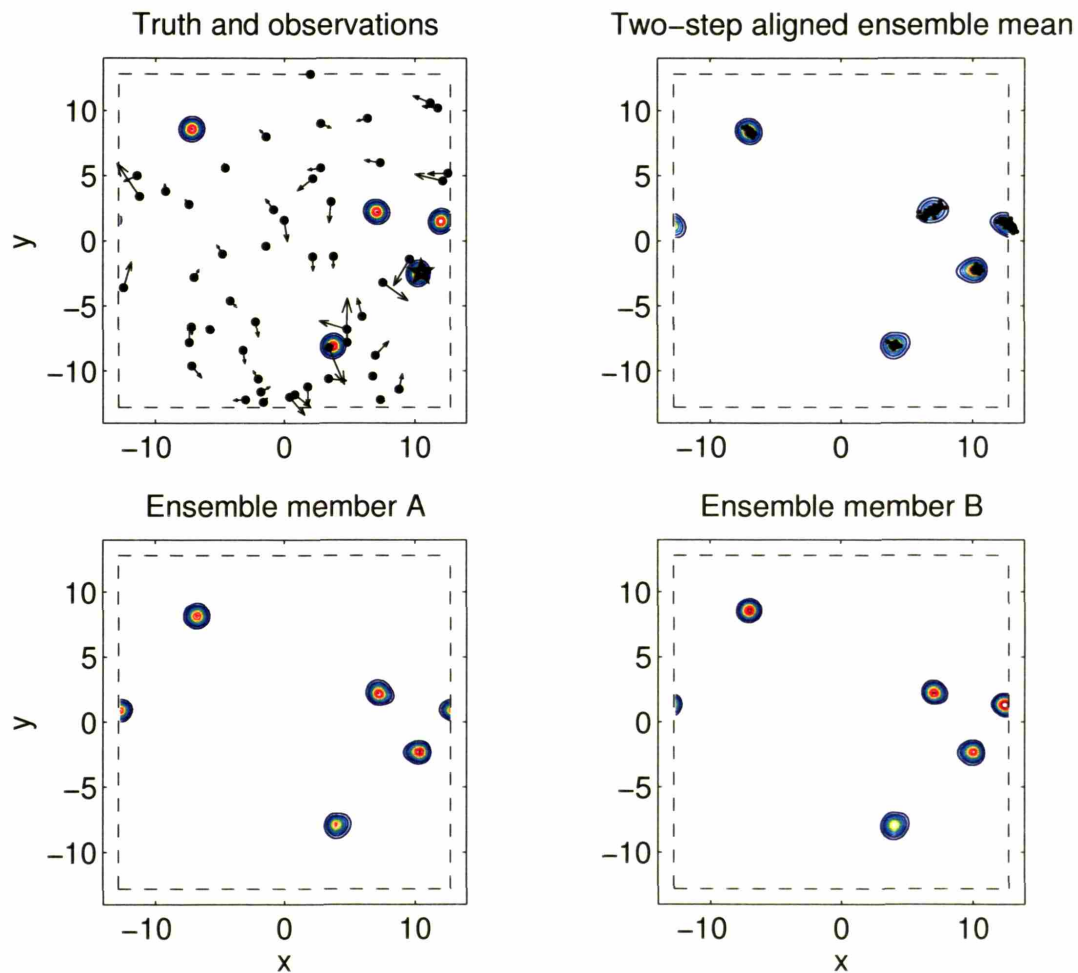


Figure 6-10: The four panels here show the same information as shown in figure 6-7, only for the two-step ensemble states after the first step of DA.

The results of the second step of DA are shown in figure 6-11. The upper left panel again shows truth and the observations for comparison. The upper right panel shows the analysis ensemble mean. The analysis mean looks only slightly different than the aligned ensemble mean in the upper right panel of figure 6-10. However, it is apparent that the second step has tried to further reposition one of the vortices in individual ensemble member A, shown in the lower left panel. The vortex near (8, 2) has been rendered oblong and there

is a small pool of negative vorticity that has emerged (one dotted black contour is present, though perhaps difficult to see). These sorts of slight deformations are to be expected as there is still some slight non-Gaussianity present, one can always encounter a particularly bad observation, and the employed alignment scheme is not ideal. Still, it is apparent that any feature deformations are not severe at this point and not likely to derail the filtering process.

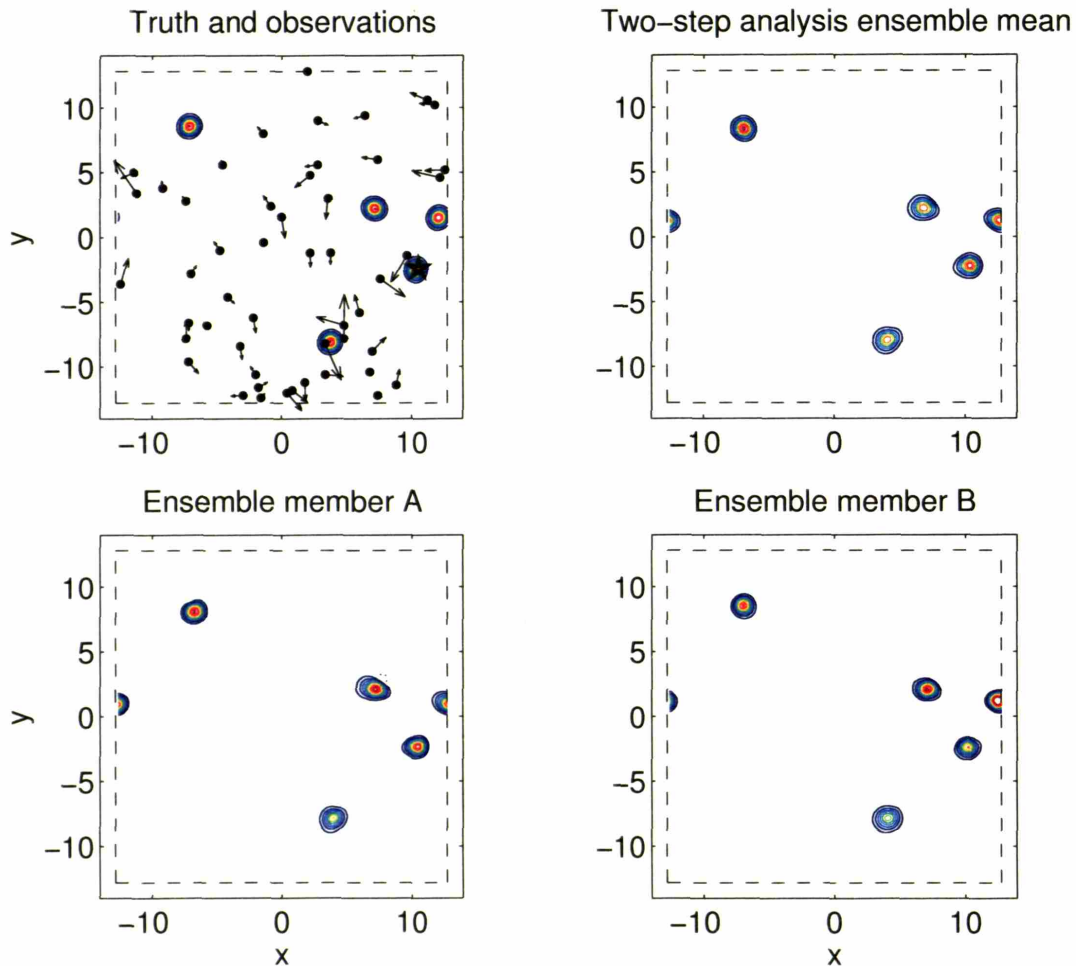


Figure 6-11: The four panels here show the same information as shown in figure 6-7, only for the two-step analysis ensemble states (i.e., after the second step of DA).

As indicated in figure 6-11, the two-step method is allowing the barotropic ensemble to effectively operate as a point vortex model. To confirm this, figure 6-12 plots the prior and analysis ensembles of the point vortex model (left panel) alongside the prior and analysis positions used in the first step of the two-step method (right panel). The prior ensembles are the open blue circles, and the analysis ensembles are the red dots. Note that the wrap-

around effect of the periodic boundaries needs to be undone for sensible DA of positions, hence some of the points are located outside the demarcated domain boundaries. It is apparent that the prior barotropic model's positions were actually less uncertain than the point vortex model's. Also, their analysis positions are very similar in spread, though they are not precisely collocated: the far right boundary vortex differs slightly. On the whole, within this observational parameter range, we conclude the two-step method's success is due precisely to its motivation of incorporating Lagrangian position information previously inaccessible to linear/Gaussian DA methods.




Figure 6-12: Both panels show prior and analysis ensemble distributions of positions in association with the fourth DA cycle of the comparative OSSE. The prior ensemble is denoted by the open blue circles, and the analysis ensemble is denoted by the red dots. The left panel shows ensembles from the point vortex model, and the right panel shows the diagnosed and analyzed positions from the first step of DA in the two-step method.

### *Tracing an individual ensemble member through a DA cycle*

In this section we show the steps involved in the barotropic model EnKF ensemble's eventual failure. Figure 6-13 traces the evolution and analysis of a particular ensemble member through three successive DA cycles. The top row shows the prior (left panel) and analysis (right panel) states for the second assimilation; the middle row shows the prior and analysis states for the third assimilation; the bottom row shows the prior and analysis states for the fourth assimilation. The bottom two panels are comparable to the prior and analysis states shown for ensemble members A and B in figures 6-7 and 6-8. We did not

choose one of those members since they have already been designated as particularly “good” and “bad,” whereas this one is more representative of “average,” and it is a particularly demonstrative example.

In each panel the relevant positions of truth have been marked with x’s, and the observed vortex further marked with a black star. The first row shows that the prior state was in fair agreement with truth (top left panel). Its update still has five vortices present, though they have become slightly deformed in shape from the EnKF’s attempts to perform fine-tuned repositioning. By the next observation time (middle left), it is apparent that the deformities have disappeared (through a combination of the numerical filtering and the circular vortex patch being a more preferred state), however, the states’ positions have fallen out of agreement. In particular, the observed vortex position disagrees by at least a full vortex diameter. The update of this state is observed to split the coherent vortices into smaller vorticity patches. The three northernmost vortices have all been severed to some degree into “double core” features, with the observed vortex completely severed. When integrated to the next observation time (bottom left), we see that, again, some of the noisiness of the previous analysis has been ameliorated, but that there are still extra and differently shaped features present. Its analysis is as poor as those shown in the bottom panels of figure 6-8. Were we to diagnose the ultimate reason for this decline in correspondence to truth (at least for this ensemble member), we would have to point to the large prior-observation misfit in the middle left panel leading to an analysis of a different character (middle right panel), that is, no longer having five well-formed features. One may also make the case for the slight vortex deformation that occurred in the analysis shown in the top right panel as to why the prior-observation misfit was so large in the middle left panel. This is the nature of cycling compounding errors through DA — the ultimate reason, of course, is that the involved error PDFs cannot validly be considered Gaussian.

As a comparison, figure 6-14 shows the same ensemble member evolve through the two-step DA scheme (the intervening aligned ensembles are not pictured). Looking at the panels, we see similar plots to those shown in figure 6-14 in the left column. Most notably, we see a large prior-observation misfit in the middle left panel. However, we see quite different plots in the right column. There are no raised issues with the EnKF trying to reposition a feature on its own as the alignment scheme is accomplishing that task (based on a DA problem that an EnKF is solving). We note that many of the vortices in the analysis panels (right column) are slightly deformed. We feel this is unfortunately an artifact of our alignment scheme (note the faint triangular shapes), but we note that it does not seem to have made much of an impact, at least yet.

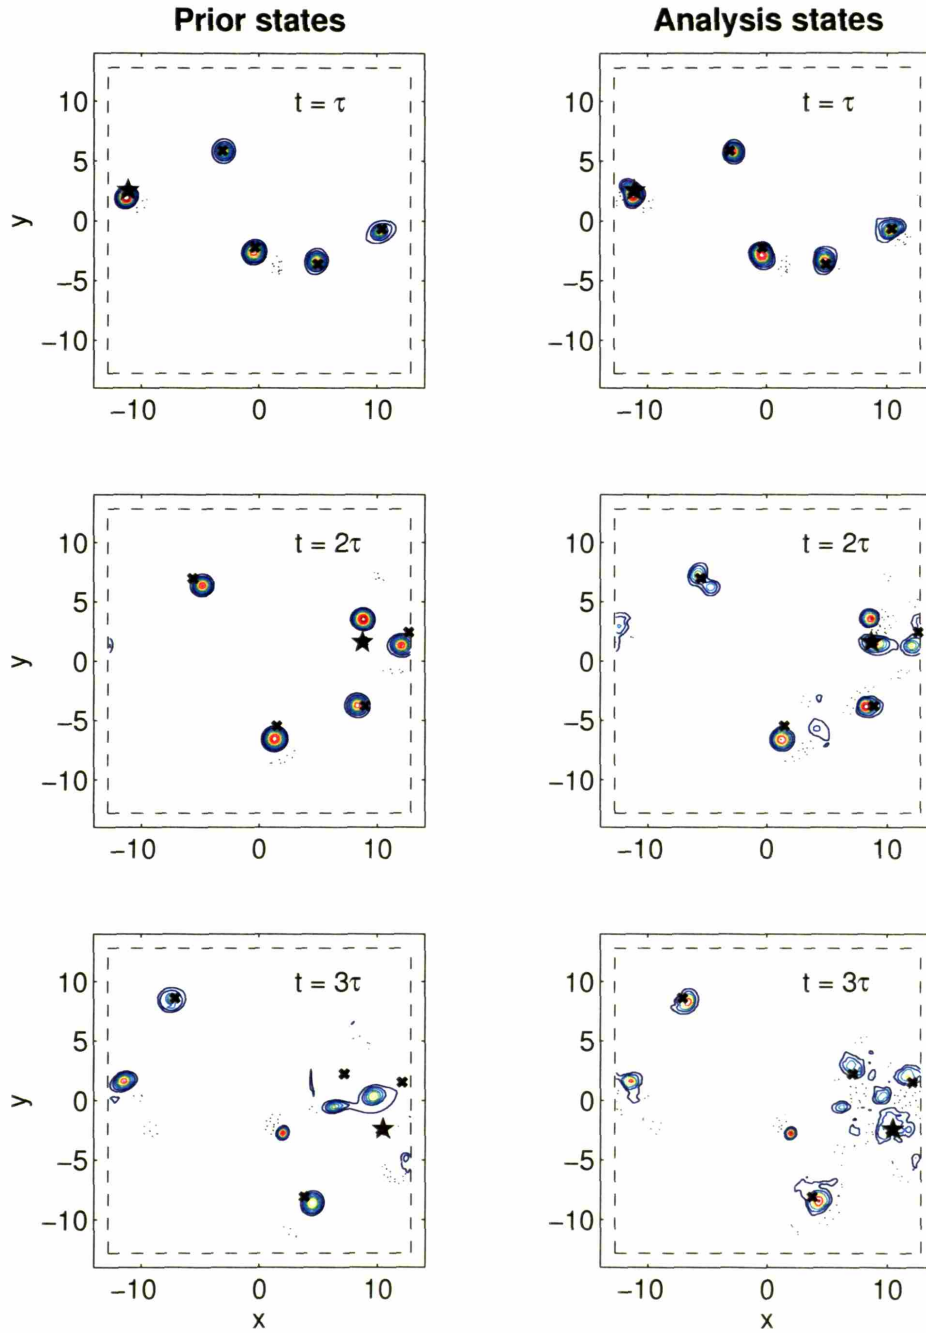


Figure 6-13: The left column of panels shows the states of a particular EnKF ensemble member just prior to DA. The right column shows that same member after DA. The top row shows the states associated with the second DA cycle, the middle row the third, and the bottom row the fourth. The contour interval and convention is the same as has been established in the preceding figures. Truth's positions are shown in each panel by the overlain black x's. The observed vortex at each time is further denoted by the black star.




Figure 6-14: This figure shows analogous information to that shown in figure 6-13, only for the two-step method. The ensemble member considered here is the same one as considered in figure 6-13 for the EnKF.

## *Discussion*

We begin by presenting a summary figure of all three ensembles, shown in figure 6-15. Each of the panels has already been shown in the preceding figures, but we include them here for easy side by side comparison. The top left panel shows the true state of the barotropic model. The top right panel shows the point vortex analysis ensemble and its mean. The ensemble members are plotted as blue dots, but are largely obscured by the other quantities plotted (a testament to their small ensemble dispersion). The ensemble mean positions are plotted as red dots. The thick black contours denote the first positive vorticity contour from the truth state shown in the upper left panel, and they are included to provide easy visual confirmation that the ensemble's positions are accurately located. The bottom two panels show the contoured analysis means of the two barotropic model ensembles, the left panel showing the EnKF ensemble, and the right panel showing the two-step ensemble. Truth's contour is again overlain to provide easy visual comparison with truth's locations. Note that all three ensembles, particularly the two barotropic model ensembles, are capturing the correct number of vortices and have them located in about the correct positions. We know from further analysis that the barotropic model EnKF ensemble is comprised mainly of states that have lost this correspondence, but their aggregate is observed to be somewhat accurate. Continued cycling of these ensembles shows that the EnKF ensemble continues to degrade (similarly to what we observed in the KdV system in figure 5-13), whereas the other ensembles can maintain an accurate correspondence with truth.

We note, however, that the two-step approach to DA is not going to be able to solve all problems. Since the first step of the two-step approach employs an EnKF, it is still subject to the assumptions of having nearly Gaussian PDFs. If one obtains significantly non-Gaussian PDFs for feature positions, then this approach is not likely to work. There is another potential brand of problems related to our chosen implementation; problems can arise with the alignment scheme. While our piecewise linear mapping scheme seems adequate for much of the features aligning we have encountered in our OSSEs, it has some limitations that become more obvious with larger numbers of features that need significant realignment. In particular, the alignment scheme does not seem equipped to handle rotating two vortices about one another more than  $70^\circ$  or so — at  $90^\circ$  piecewise linear mapping will often map a triangular tile into a collapsed line, thus removing some of the intervening gridpoints (though algorithmically, this resolution is regained elsewhere). This is a consequence of the analysis positions violating the prior positions' triangulation (i.e., triangulation lines intersect one another in the analysis position configuration). This is an idiosyncratic deficiency of the specific alignment scheme employed, and as we are aware that more robust schemes are

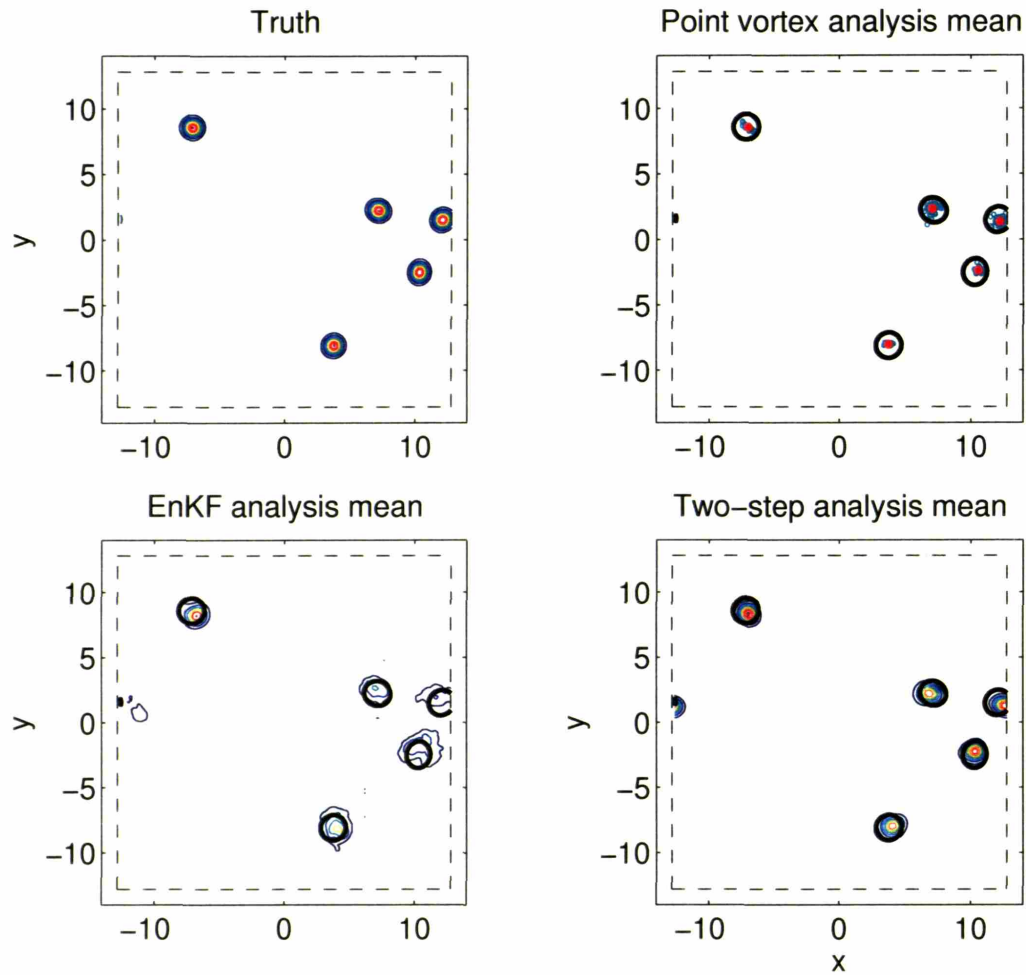


Figure 6-15: The top left panel shows truth at the 4th assimilation cycle. The top right panel shows the point vortex analysis ensemble members (blue dots) and their mean (red dots). The thick black contours correspond to the first positive vorticity contour of truth. The bottom left panel shows the barotropic model EnKF analysis ensemble mean, again with truth’s contours overlain, and the bottom right shows the two-step method’s analysis ensemble mean. The contour interval and convention are the same as in the preceding figures.

possible, we feel that this brand of problems arising from the two-step ensembles should not dwelled upon too heavily.

We should view the observed improvement of the two-step approach over the standard EnKF as very encouraging. In this point vortex scenario, the core dynamics are *de facto* Lagrangian, and as such this scenario offers an ideal setting for the two-step approach to showcase its advantage. *This is a setting where we know the normal EnKF will not work (see sections 4.2.4 and 5.3.3), yet the two-step method is able to succeed using only*

*traditional linear/Gaussian DA and working with the same amount of information obtained external from the ensemble; its very success comes from its utilization of information contained within the ensemble.* In practice, it would seem the two-step method’s success rests largely on the ability of the alignment scheme to operate with minimal distortion, and, fortunately, we find in most circumstances that our chosen scheme is adequate for this purpose.

It should be noted that the relative successes of the filters depend on the Gaussian vortex size parameter  $d$ . When  $d$  is small enough so that the vortices remain separated enough to appear to one another as point vortices (as in the results shown above), then the barotropic model is able to reproduce the evolution of the point vortex model, and we find that the barotropic model EnKF ensemble generally fails while the two-step ensemble is able to retain good correspondence to truth. As  $d$  gets larger, the merits of the two-step relative to the EnKF are diminished. When  $d$  is large enough for the Gaussian vortices to no longer appear to one another as point vortices (e.g.,  $d = 12\Delta x$  for  $128 \times 128$  resolution), then the barotropic model cannot reproduce the point vortex model’s evolution, and we find that the barotropic model’s two ensembles become essentially equivalent for most observational parameter ranges. This is because the vorticity patches are so broad that the contribution of position errors are minimized into the acceptable Gaussian range required by the EnKF.

The point vortex example shown, while illuminating, can be seen as somewhat contrived. We next turn our attention to a more realistic example, the state estimation and prediction of a jet losing structure to barotropic instabilities.

### 6.3.3 A barotropically unstable jet

In our introduction to features and position errors, section 4.1.1, we discussed two classes of feature emergence studies that have been examined in the past: features from rotating, two-dimensional turbulence, and features as growing instabilities. These two types of studies find feature emergence from somewhat opposite approaches: the former branch studies the emergence of coherent vortices from highly randomized turbulent flows with no discernible initial structure, whereas the latter branch watches coherent vortex patches form as an initially well-structured jet deteriorates due to growing instabilities. From a predictability standpoint, either scenario would be an interesting test bed once the features have emerged, but only the latter seems practicable to examine while the features are emerging. It is for this reason we study the decay of a barotropically unstable jet.

## Model behavior

The example we choose to study is adapted from examples treated thoroughly by Flierl et al. (1987). The authors were concerned with exploring the parameter space of a barotropic model very much like ours when initialized with a barotropically unstable jet profile on a  $\beta$ -plane. They mainly varied the value of  $\beta$ , the ambient meridional vorticity gradient, and the zonal wavenumber of the initial sinusoidal-in- $x$  perturbation to the jet. They found a rich spectrum of nonlinear regimes, including what they term “vortex streets,” where the jet breaks up into a finite number of regularly spaced vorticity patches. They additionally found somewhat erratic vortex generation from using a spectrum of wavenumbers to initially perturb the jet. An example of this erratic evolution where  $\beta = 0$  is presented in figure 6-16, adapted from Flierl et al.’s figure 18. The eight panels show the vorticity profile of the jet through time as it breaks down into coherent vorticity patches, that is, features. It is this experiment that we seek to use in a predictability context: we will generate an ensemble of states each with slightly different initial jet perturbations. We choose to keep  $\beta = 0$  — a non-zero  $\beta$  tends to confine the emerging features to the original region of the jet, whereas we are seeking a scenario with fully evolving and interacting features.

Though we have resisted defining what constitutes a feature throughout this thesis, figure 6-16 gives an opportunity for us to identify a few features so the reader gains a qualitative view of what we mean. We would say that neither of the top two panels have features in them (or rather, both have a single stark feature to which we will not attempt to ascribe a single position). The third panel ( $t = 2\tau$ ) may or may not be seen as containing a feature, depending on the question of interest. By this time perturbations on the jet have begun to grow, and a wavenumber four pattern has started to emerge. One can see that there are several defined minima and maxima of vorticity. Depending on the setting, we may be tempted to designate the trough extending farthest southward (the first vorticity minimum) as a feature. In the fourth panel ( $t = 3\tau$ ), we think there is little ambiguity in identifying eight separate features, though we note that the success of the two-step does not rely on identifying all eight (there is further discussion on this below).

We generally follow the model configuration and state initialization used by Flierl et al., including the form of the numerical hyperdiffusion,  $\mathcal{N} = -\nu\nabla^4q$  (see equation (6.16); also see appendix D). The exponential wavenumber filter used in the previous subsection (see appendix D) does not affect the larger scales, while the Flierl et al. hyperdiffusive form affects all scales of  $q$ . The biharmonic term is preferable in this case as it more quickly eliminates small-scale filamentary structures attached to the emerging features, leaving relatively well-formed vorticity patches closely approximating coherent vortices. To accomplish the same

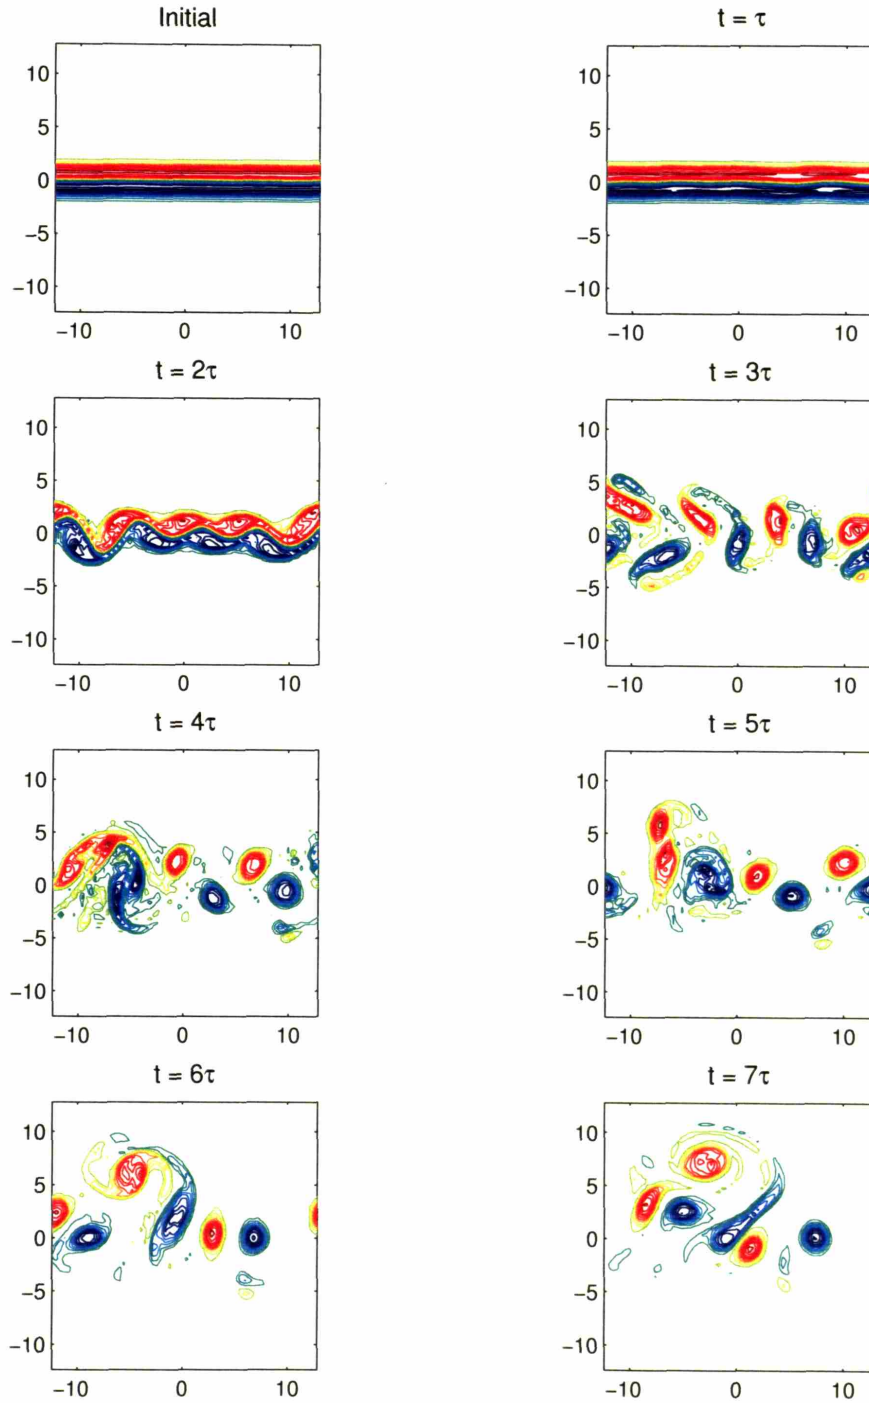


Figure 6-16: Typical erratic evolution of an initial multiple scale perturbation to an unstable jet profile. The axes are the  $x$  and  $y$  spatial directions, the field plotted is vorticity, and all panels have the same contour interval with the zero contour omitted. The panels are separated by a time interval  $\tau$  which is the observation time used in the following OSSE.

filament dissipation with the wavenumber filter requires some tuning of its threshold value, which we found difficult and did not perform.

### The jet's meridional profile

The upper left panel of figure 6-16 shows the initially perturbed jet profile, though the perturbations are of such small amplitude that they are hardly discernible. As the boundaries are periodic, any non-zero meridional jet profile with  $\beta = 0$  satisfies the necessary condition to be barotropically unstable because  $\frac{\partial^2 \bar{u}}{\partial y^2}$  will have to change sign within the domain to satisfy periodicity ( $\bar{u}$  is the zonal mean zonal velocity profile). Also, doubly periodic boundaries mean that wavenumber zero (i.e., the basin wide average) is arbitrary and constant; it is customarily assumed to be zero. The profile chosen by Flierl et al. is a relatively narrow Gaussian jet profile, implying:

$$\psi_{\text{jet}}(y) = -\text{erf}(y) + \frac{2y}{L} \quad (6.19)$$

$$u_{\text{jet}}(y) = \frac{2}{\sqrt{\pi}} \exp(-y^2) - \frac{2}{L} \quad (6.20)$$

$$q_{\text{jet}}(y) = \frac{4}{\sqrt{\pi}} y \exp(-y^2), \quad (6.21)$$

where  $\text{erf}(y)$  is the standard error function, and  $L$  is the meridional domain length.  $L$  is presumed much greater than the jet's half-width,  $\sqrt{\log(2)}$ , to ensure the jet is "narrow." These cross-sectional profiles are shown in figure 6-17. Note each profile has zero mean. The arrows in the middle panel indicate the weak basin-wide return flow in the mean velocity profile.

### Initial jet perturbations

Model states are made to go unstable, like the state shown in figure 6-16, by adding a perturbation to the jet. Flierl et al. considered so-called sinuous perturbations to the jet, sinusoidal in the zonal direction and contained within the meridional extent of the jet. To achieve the erratic behavior shown in figure 6-16, they considered a summation of such perturbations over many wavenumbers:

$$\psi'(x, y) = \delta \exp(-y^2) \sum_{n=1}^{N_k} \frac{1}{n^2} \sin(nk_0 x + \phi_n), \quad (6.22)$$




Figure 6-17: These show the meridional profiles of initial jet structure. The top panel shows the streamfunction, the middle panel shows the zonal wind, and the bottom panel shows the vorticity. The arrows in the middle panel indicate the direction of the zonal wind, showing a basin-wide weak return flow outside of the narrow jet region.

where  $\delta$  is a perturbation amplitude,  $N_k$  is the number of contributing modes to the perturbation,  $\phi_n$  are random phase factors, and  $k_0$  is a zonal wavenumber for the perturbation, assumed to be quantized within the domain length (i.e.,  $k_0 = 2\pi n_0/L$ , where  $n_0$  is the number of waves in the domain). Following Flierl et al., we choose  $n_0 = 1$  and  $N_k$  equal to the number of gridpoints, generally taken to be 64, though some cases with 128 were considered. Flierl et al. considered  $\phi_n$  to be from a uniform PDF over the interval  $[0, 2\pi]$ . We chose to follow this suggestion. This means that most ensemble members will share qualitatively similar evolution (i.e., the number of wave forms to develop and cut off), but they will potentially share no zonal position correspondence (i.e., their relative phasing will be uniformly random). By the arguments and examples already presented in this thesis,

this will clearly disrupt traditional DA if large position discrepancies are allowed to develop. Therefore, we need to be conscientious during the features' incipient stages — it is conceivable (though unlikely) that DA applied to the ensemble when the states still resemble the top right panel of figure 6-16 will be sufficient for locking in the members' relative phasing.

### Experimental set-up

Following Flierl et al., and as in section 6.3.2 with the barotropic point vortices, we make no attempt to intrinsically configure the model with geophysical scales, though we do perform DA knowingly in reference to typical frequencies used in NWP by way of the model's error doubling time. In the experiment that follows, the error growth regime was purposefully placed in a range where the two-step method is successful and traditional ensemble-based methods are not. As examined throughout chapter 3 and in the experiments shown in chapters 4 and 5, the error growth regime is largely controlled by the observation network. In particular, to mimic NWP scenarios, the observation time (the period between observations) ought to be set in relation to the system's error doubling time. As discussed in section 4.2.3, typical synoptic-scale NWP seems to work well for an observation time around one eighth to one quarter the error doubling time; however, this is largely excluding the influences from mispositioned coherent features. The error doubling associated with feature growth can be quite rapid, rapid enough to violate the tangent linear approximation within an observation time (e.g., see section 5.5.2). Hence, for the purposes of DA we assume the system in this configuration has the same error doubling time as the point vortex model configuration considered in the previous subsection, though we suspect that this is an *overestimate* of the actual error doubling associated with the jet going unstable. This overestimation is acceptable since performing DA at this frequency will place the system in a regime of appreciable nonlinear error growth where the traditional ensemble-based filtering methods fail. Also, this scenario will give an opportunity for the two-step method to demonstrate its ability to correct the non-Gaussian PDFs associated with the dispersing features. Hence, we choose an observation time of 12 hours, where the notion of "hour" and "day" come from equating 480 model  $\Delta t$ 's (the error doubling time for  $64 \times 64$  resolution) to two atmosphere days.

We use the following values for our model integrations, where time and distance units have been nondimensionalized by the maximum velocity scale and the jet's half-width (i.e., "model distance" and "model time"). The model run parameters are taken directly from Flierl et al.'s study. We use a square model domain with side length  $L = 25.6$  centered at the origin. We use  $64 \times 64$  spatial resolution (i.e.,  $n = 4096$ ), a timestep of  $\Delta t = 0.0664$

(in model time), an ambient vorticity gradient  $\beta = 0$ , and a hyperdiffusion coefficient of  $\nu = 6.64 \times 10^{-5}$  (in units of  $[\text{dist}^4 \text{time}^{-1}]$ ).

For our ensemble generation, we again generally following the perturbation strategy of Flierl et al. In reference to equation (6.22), we use: a perturbation amplitude  $\delta = 0.02$ , a fundamental perturbation wavenumber  $k_o = 2\pi/L$  (meaning  $n_o = 1$ ), and the number of contributing perturbation wavenumbers  $N_k = 64$ . For each of  $N = 50$  ensemble members, we generate a uniformly random phase perturbation over the interval  $[0, 2\pi]$ .

For observations, we randomly select 25 fixed “station” observation locations within the domain prior to beginning the OSSE. Since each station measures a two-dimensional vector quantity, there are 50 observations taken. This amounts to about 1% of the size of the state vector, fairly sparse coverage. For simplicity in the observation operator, the observation locations are chosen to be at model gridpoints. Unlike in the barotropic model experiments run in section 3.4.2 where we directly observed vorticity, here we observe only the fluid velocity, or “wind,” at the stations. This is the same observation strategy we used in section 6.3.2 with the barotropic point vortex experiments.

The error structure for wind observations can either be specified by magnitude and angle or component-wise, and here we choose component-wise errors, though the alternative case is easily implemented. We choose to prescribe an absolute error scale for the wind of 5% of the jet’s initial maximum velocity (nondimensionalized to be unity, so our velocity error scale is just 0.05), rather than a proportional error scale (i.e., 5% of the actual wind magnitude measured versus 5% of fixed scale). This is done mainly out of convenience so that the observation error covariance matrix is fixed rather than dependent on the value of the observations, but it does raise an interesting issue about the relative value of observations. The largest “signal” in velocity is associated with the jet core and later the coherent features — if observations have an absolute error scale, then stations observing in relatively quiescent regions do not provide as much information to the estimation process as those in active regions based on their signal to noise ratios; if observations have a relative error, then all stations provide the same amount of information on a signal to noise basis, but the stations sampling the most interesting regions are consequently trusted the least. Probably the most realistic approach is to assume a mixed error model (see section 5.2) for observational error that is partly additive and partly multiplicative, but we resist doing so in these experiments as they are already heavily idealized to the point where such improvements may seem inconsistent and it could make our results more difficult to interpret.

In addition to the fixed observing network, there is one position observation available at each observation time, even though many features may be identifiable in the model domain

at any given time. Once identifiable features have emerged (we address the meaning of “identifiable” below), one feature is observed, though not necessarily the same feature each observation time. We assume an observational uncertainty in position of one gridbox, based on the idea that position should be easy to place from, say, remote sensing data though difficult to pinpoint exactly. This of course drives right at the heart of identifying features. We have experimented with possible schemes to act as our feature identifier, that is, our  $\mathcal{F}[\cdot]$  operator (see section 6.1), and though we have found plausible candidates (in the same spirit as our drive to find a plausible alignment scheme), none are robust enough to handle the irregular and erratic feature development encountered in our jet experiments. Devising an  $\mathcal{F}$  operator is easier for the point vortex scenario since we know *a priori* how many features to expect and their expected structure, but when a jet is going unstable, neither piece of information is known. This thesis is not about how to effectively identify features, rather it addresses the power of ensemble-based methods augmented by such information *given that it can be found*. Therefore, we eventually settled for an inelegant  $\mathcal{F}$  that was certain to work but very difficult to make objective, namely the process of manual identification. In general, we identified featuredness by several concentric closed contours. However, as discussed in sections 4.1.2 and 6.2.3, there are atmospheric scientists working on more sophisticated, automated methods for the problem of feature identification. Also, this is certainly an issue for mesoscale meteorologists, particularly those concerned with forecast verification (e.g., Nachamkin 2004).

As in all systems studied in this thesis, the barotropic jet system can of course be observed densely and frequently enough so that both methods work, but we omit figures from those experiments. We instead purposefully choose a situation where traditional ensemble-based methods fail, and this makes it difficult to compare the two-step method and an EnKF beginning from ensemble initialization because the EnKF quickly loses correspondence to truth and diverges. There are two approaches one can take to combat this, either to observe often and accurately until the features have emerged and then change the observation frequency, coverage, and/or accuracy, or to use only the two-step method for several DA cycles until the features have emerged and then update the same evolved ensemble with both DA systems. Here we have opted for the second method. This means that the figures we show will be of a more kinematic nature. Neither of these methods is ideal, but we proceed since this experiment is mainly for the purpose of demonstration rather than gathering verification statistics or ensemble diagnostics (as in chapter 3).

## OSSE results

Here we present one full ensemble update comparison where both an EnKF and the two-step method are applied to the same initial ensemble distribution. This initial ensemble will also be termed the “forecast” ensemble since it was taken from an OSSE.

### *Initial ensemble states*

Figure 6-18 shows the truth state, the observations, and some prior ensemble information at the fifth observation time (roughly comparable to the sixth panel in figure 6-16). The contours in the upper left panel show the truth state  $q$  with the zero contour omitted (red is positive and blue is negative). The black circles show the fixed observation stations, and the vector arrows denote the wind observations. The black star over the far left red vorticity patch shows the position observation for this DA cycle. The upper right panel shows the initial ensemble mean with the same contour interval as truth’s in the left panel. It is evident that the features’ relative dispersion across the ensemble has smoothed the shapes and lessened the amplitudes of the features shown in the mean. To give the reader a sense of the relative dispersion among the ensemble members, the bottom two panels each show a particular ensemble member. The members are not extreme members in any obvious sense, rather they were chosen to represent the typical spread of the features. Note that they are not too dissimilar, for example, they have the same number of features with comparable sizes, shapes, and strengths. This means that drawing correspondence between features across the ensemble is well-posed. The ensemble’s internal correspondence and likeness to truth is a testament to the four applications of the two-step method that were needed to obtain this ensemble.

### *EnKF results*

Figure 6-19 shows the corresponding updates from using a traditional EnKF scheme. The upper left panel repeats truth and the available observations from figure 6-18 for reference. In order to execute the EnKF with use of the position observation, we must apply the  $\mathcal{F}$  operator to each member to diagnose the position of the observed feature within that member so that an innovation can be evaluated.<sup>6</sup> The upper right panel of figure 6-19 shows the analysis ensemble mean from the EnKF, and the bottom two panels show the respective updates of the particular forecast members shown in the bottom two panels of figure 6-18. The analysis mean is reasonable; indeed, it is an approximation of the minimum

---

<sup>6</sup>The top panel of figure 6-20 indicates the ensemble spread of the observed feature’s position (one dot for each of the 50 ensemble members), though that figure will be discussed at length below

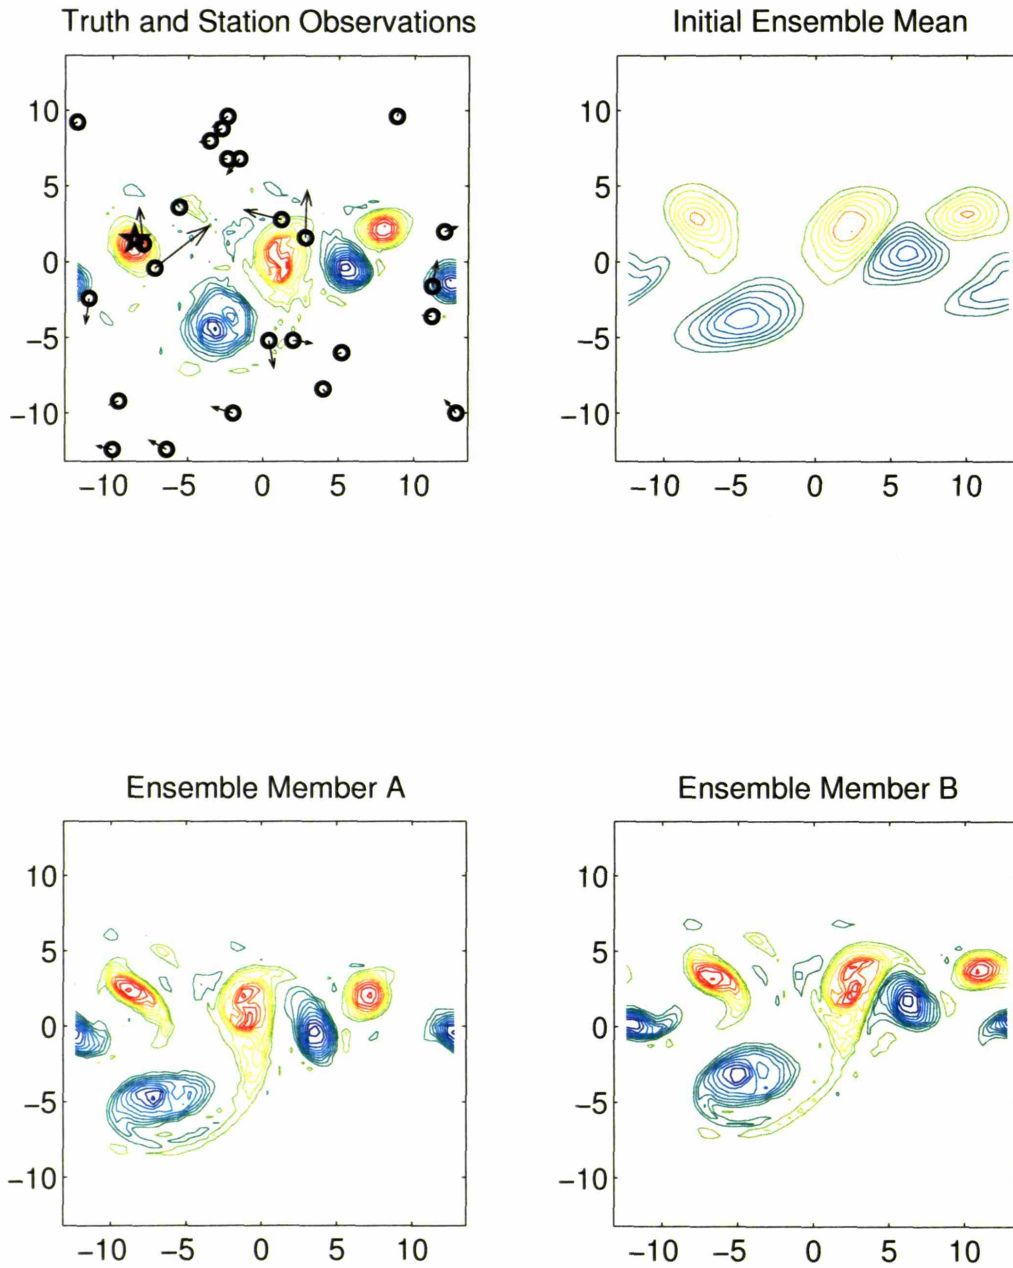


Figure 6-18: The contours in the upper left panel show truth's  $q$  field, the black circles show the fixed observation station locations, the vector arrows show the wind observations, and the black star shows the position observation. The upper right panel shows the forecast ensemble mean with the same contour interval as in the upper left panel. The bottom two panels each show a representative individual ensemble member.

variance estimate based on the information at hand. It has roughly the right number of distinct vorticity patches in roughly the correct positions. As it is a mean, the features

are still rather smoothed and weakened, though they have been tightened over the prior mean shown in figure 6-18. There is a banded structure (i.e., positive-negative-positive) associated with the far left feature whose position was observed. This is due to the limited ability of an EnKF to reposition features (see 5.3.3). As the banded structure is imprinted in the ensemble mean, we expect it is in most of the individual members, and the bottom two panels confirm this. In addition to the banded structure, the negative vorticity patch at around (5,0) has been updated into two smaller patches in ensemble member A (lower left panel). Also, the large feature near the origin is quite weak and “noisy” (i.e., contains many weak, poorly organized local maxima in vorticity). Ensemble member B (lower right panel) also has the banded structure and exhibits noisy vorticity patches. It is apparent by eye that the truth state in the upper left panel is distinguishable from the ensemble members that comprise the mean. This is all analogous to the soliton case shown in the top panel of figure 5-11.

#### *Two-step method results*

When we implement the two-step method, we can take further advantage of features within the ensemble that are not necessarily observed. Though there are six identifiable features in each ensemble member, we only focused our  $\mathcal{F}$  operator on finding four of them. It is important to stress that the two-step method accepts whatever position information the user cares to give it: one need not diagnose every feature’s position. A feature’s position here is the geographical location of its maximum value in  $q$ , though this raises an interesting issue as the position of the maximum in  $q$  need not be collocated with the position of the minimum in  $\psi$  or the maximum (or minimum) in the wind magnitude,  $|\mathbf{V}|$ . What is the proper diagnosis of position? In the case of the BVE, position is probably best diagnosed by vorticity since it is the master variable and can be treated piece-wise linearly (i.e., it is sensible to treat an individual piece of the vorticity field), but it is not so clear in general. In practice, position ought to be diagnosed by averaging together the estimates derived from several different fields, and the differences between them ought to further inform the expected uncertainty in the prior estimate of positions. Hoffman et al. (1995) address this issue by suggesting that variables be transformed to forms convenient for rearranging (e.g., potential vorticity or potential temperature) before analyzing them as such. By working with our idealized model, we avoid much of this messiness, and this is acceptable since this example is not meant to be as realistic as possible, only a proof-of-concept that the two-step method can successfully be extended to multiple dimensions.

The top panel of figure 6-20 shows the forecast ensemble mean with black dots denoting

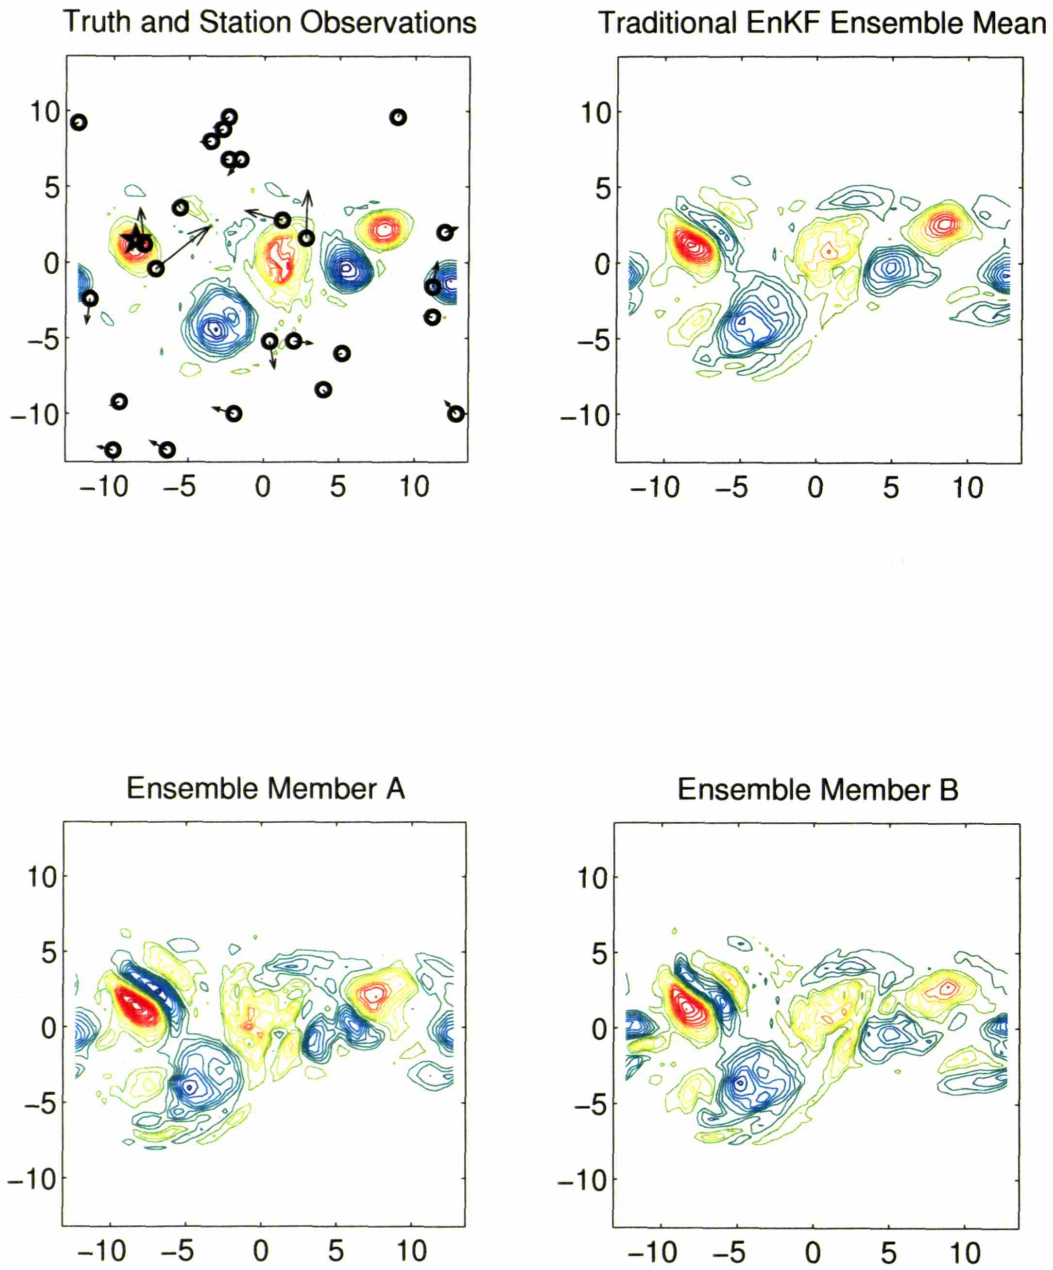


Figure 6-19: The upper left panel repeats the upper left panel of figure 6-18 and is mainly included for reference. The upper right panel shows the EnKF analysis ensemble mean with the same contour interval as that used in the first panel. The bottom two panels show the individual updates of the initial members shown in the bottom two panels of figure 6-18, again with the same contour interval.


the diagnosed positions of the four central features (including the one that was observed in truth). This information is augmented to the ensemble state vectors, and the augmented

ensemble is updated by an EnKF. The analysis feature positions,  $\mathbf{p}_j^a$ , are plotted as black dots in the bottom panel of figure 6-20. The dots are overlain upon the forecasting ensemble mean of  $q$  to emphasize that after the first step of the two-step method, it is only the  $\mathbf{p}_j^a$  component of the analysis (augmented) state vectors that are of interest. The position dispersion has been tightened quite a bit. Since all available observations have been used in this first step, it is not immediately clear what aspect of the update leads to the great decrease in uncertainty, though it is clear that the position observation helps update the observed feature's position. By recomputing the ensemble updates with a selective batch of observations, we have determined that the other three feature positions are updated so accurately because of i) covariance information that exists between the four diagnosed feature positions, and ii) the available fixed station observations, with the latter having a stronger effect than the former.

Applying the alignment scheme to the ensemble members, we obtain the results shown in figure 6-21. The upper left panel again shows truth and the observations. The upper right panel shows the aligned ensemble mean, comprised of members whose bases have been remapped so that their features are nearly collocated. The contours of this panel should be compared to those in the bottom panel of figure 6-20 to see how the mean has been tightened up. Note, however, that the two features not included in the alignment process have not tightened as noticeably as the other four. The bottom two panels show the aligned individual ensemble members. Comparing these to the bottom two panels of figure 6-18, we see that the alignment scheme has moved the large negative vorticity patch in the third quadrant of ensemble member A eastward while leaving the other features relatively untouched, whereas many of the vorticity patches in the northern half of ensemble member B have been shifted southwestward. There has obviously been some distortion to the original members, but quite like Hoffman et al. have come to re-term "distortion" as "adjustment" in the evolution of the feature calibration and alignment (FCA) literature (see section 4.1.4), we submit that the alignments shown in figure 6-21 do not deserve the negative connotations of "distortion." Indeed, the alignment scheme is simply enforcing the results of a state estimation procedure. Also, the states shown in the bottom two panels of figure 6-21 are less distorted (by almost any sensible measure of distortion one can concoct) than those in the bottom two panels of figure 6-19, and those were generated by applying traditional ensemble-based DA.

With the aligned ensemble in hand, the two-step method completes its analysis by applying a traditional EnKF to the aligned ensemble. The results are shown in figure 6-22. The upper left panel again repeats the truth and observations for comparison. The upper

Initial Ensemble Mean with Initial Feature Positions



Initial Ensemble Mean with Analyzed Feature Positions




Figure 6-20: The top panel shows the forecast ensemble mean field of  $q$  with the same contour interval as has been used in the previous two figures. The overlain black dots show the diagnosed ensemble positions of the four central features. The bottom panel shows the same ensemble mean  $q$  field as in the top panel to emphasize that the first step of the two-step method does not change the state vector values, only the diagnosed position values. The analysis feature positions are shown by the tightly clustered black dots.

right panel shows the two-step analysis ensemble mean. The analysis mean is improved even over the aligned ensemble mean: the two features not included in the alignment process

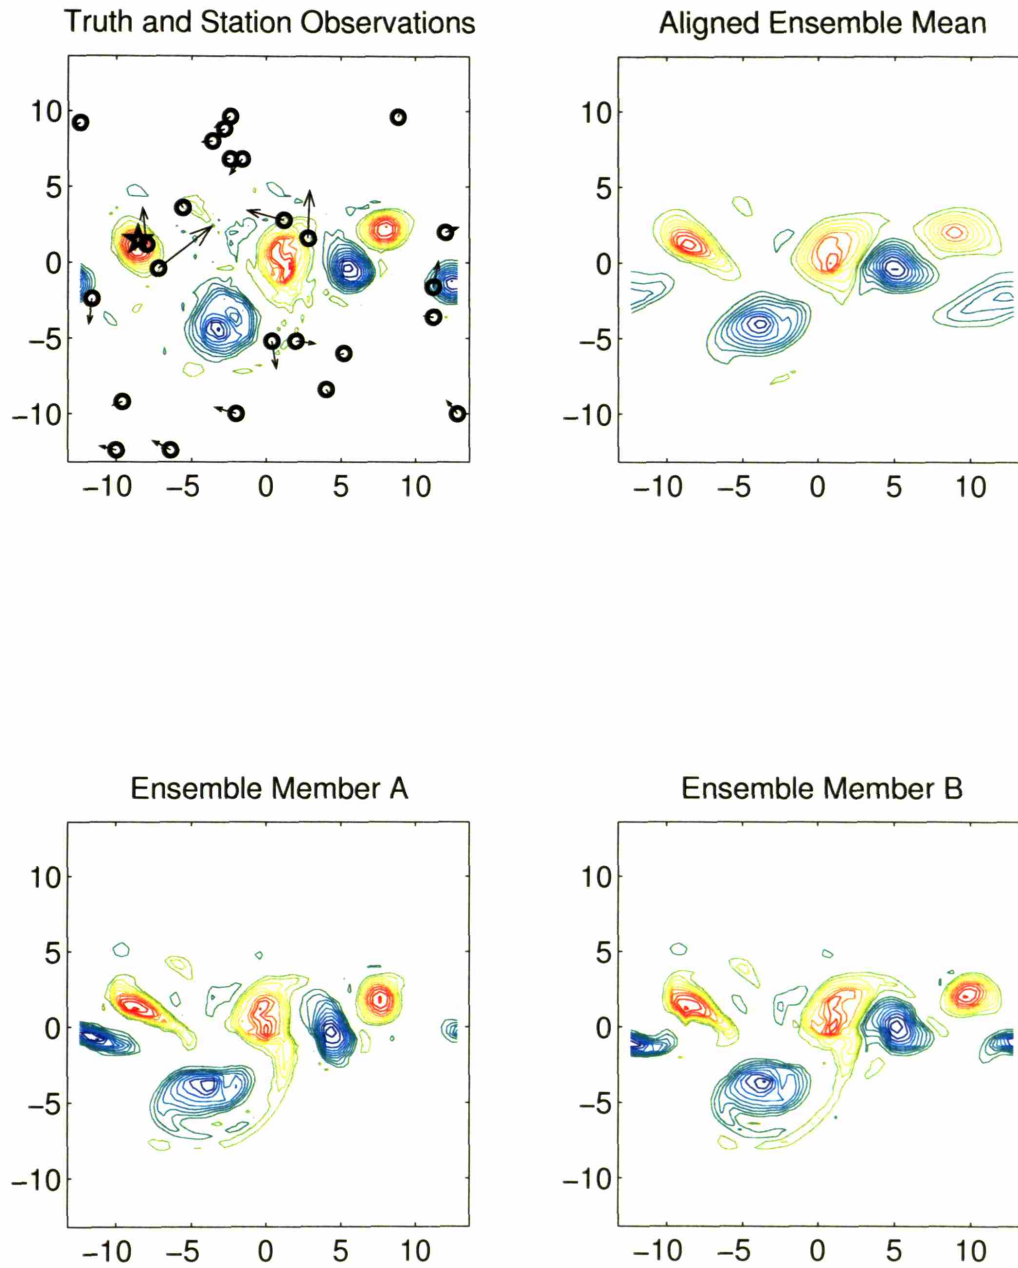


Figure 6-21: The upper left panel repeats the upper left panel of figure 6-18 and is mainly included for reference. The upper right panel shows the aligned analysis ensemble mean with the same contour interval as that used in the first panel. The bottom two panels show the aligned individual ensemble members, again with the same contour interval.

have been further tightened. The individual members, however, have incurred some of the noisiness observed in the traditional EnKF update. Also, the upper right positive vorticity patch in ensemble member B has a small negative vorticity nodule to its north

now, consistent with the repositioning “wakes” considered in section 5.3.3. This is one of the two features not included in the alignment process, and one can see in the lower right panel of figure 6-21 that the second step of the two-step has tried to move the patch southward. As we saw in the soliton example in chapter 5 (see bottom panel of figure 5-7), the second step’s success is still conditioned on the first step and the alignment scheme removing most of the non-Gaussianity in the forecast ensemble. Here we conclude that the first step was able to remove much of the non-Gaussianity, but not all of it.

### *Discussion*

The updates from the two-step method show a great improvement over the EnKF updates shown in figure 6-19, particularly for the individual updates in the bottom two panels. In fact, it is mostly in regard to the individual updates that ground has been gained. The two-step method has helped maintain an ensemble of plausible indistinguishable states when compared to truth. As ensemble-based DA seeks to approximate the minimum variance estimate, we find that the two methods produce comparable means. To make their comparison easier, figure 6-23 shows truth and the two analysis ensemble means. Overlain on the analyses are black contours from truth’s  $q$  profile (like a spaghetti plot, see section 4.2.4) to allow easy comparison of feature positions. It is evident that the obtained ensemble means from the two different methods both reasonably capture the structure of truth, that is, they both have representations of features with the correct signs of circulation in about the correct positions. The EnKF ensemble analysis mean has some “spurious” vorticity structures that have arisen as a product of the DA (most notably the banded structure associated with the upper left positive vorticity patch), but as a succinct summary of the PDF representing our best estimate of truth, probably either is acceptable. However, as an ensemble of states that are statistically indistinguishable from truth and fit to serve as initial conditions for the next batch of integration, the two-step analysis ensemble is the only of the two that is acceptable.

Here we echo the conclusions drawn from the above point vortex experiments: we should view the observed improvement of the two-step approach over the standard EnKF as very encouraging. Even with a less-than-perfect alignment scheme, we find we are able to extend the observational parameter range over the traditional (“state-of-the-art”) DA methods in a scenario approaching the behavior of real weather. As discussed in section 4.1, we do not ultimately expect the two-step method to be necessary in large-scale synoptic weather forecasting since the position errors associated with synoptic waves are typically much less than their relative length scales (see section 5.2.3), but nevertheless, this example shows

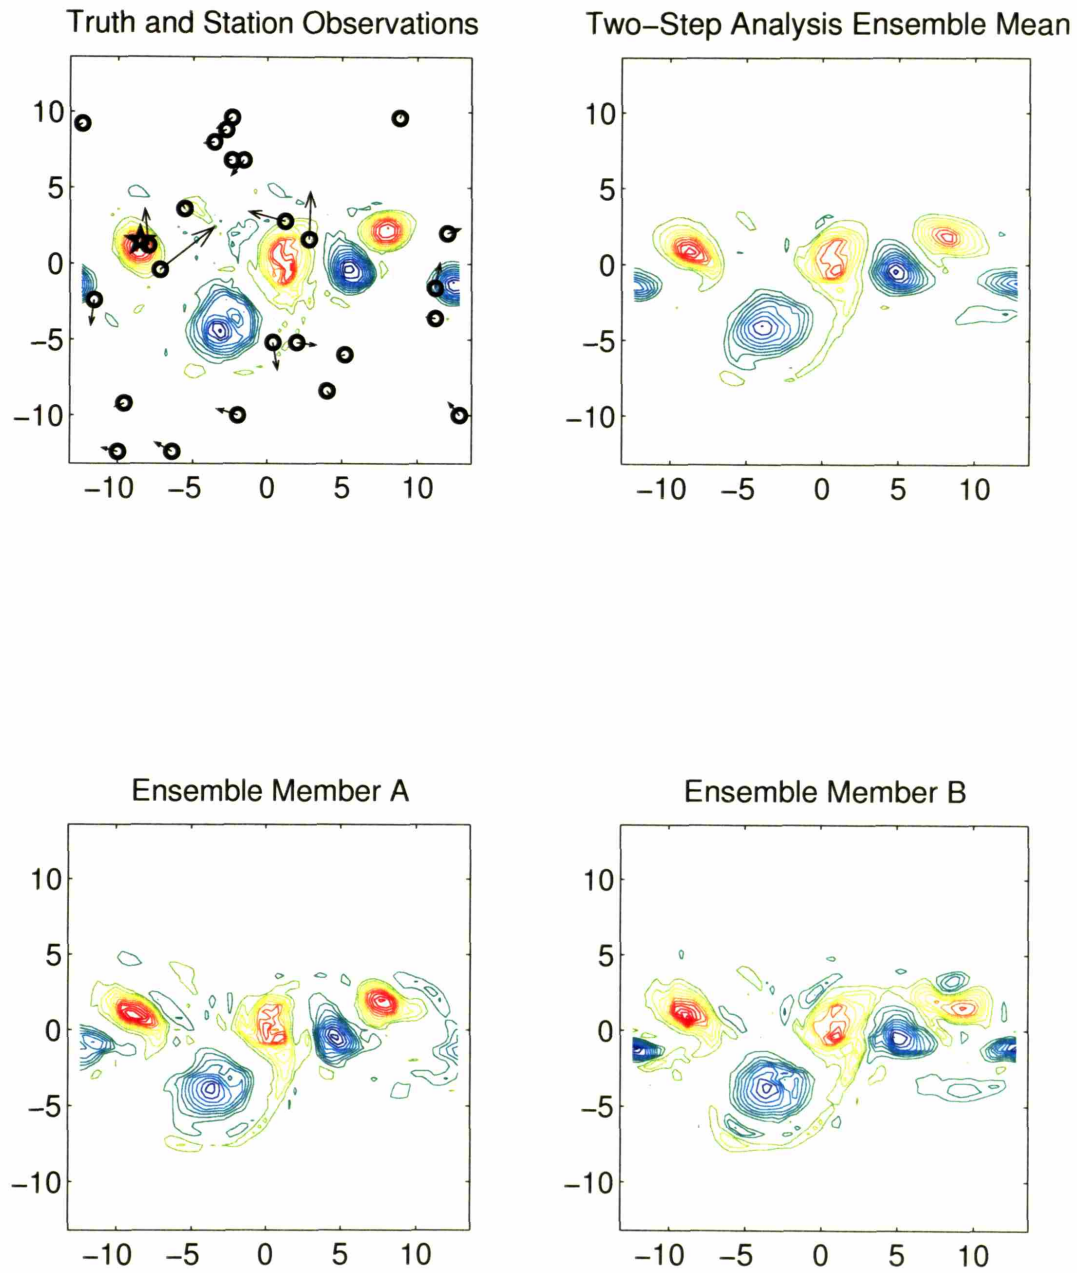


Figure 6-22: The upper left panel repeats the upper left panel of figure 6-18 and is mainly included for reference. The upper right panel shows the final two-step analysis ensemble mean with the same contour interval as that used in the first panel. The bottom two panels show the individual analyzed ensemble members, again with the same contour interval.

that multiple, irregularly shaped, interacting features should not be a problem for a well-designed operational two-step method.




Figure 6-23: The top panel shows the truth state, the middle panel shows the EnKF analysis ensemble mean, and the bottom panel shows the two-step method analysis ensemble mean; all panels have the same contour interval. The overlain black contours in the bottom panels are from a selected value of truth's  $|q|$  field, and they are included to make position comparison easier.

## 6.4 Summary

In this chapter we have successfully extended the two-step method of data assimilation developed in chapter 5 to spatially extended fields in two dimensions and with multiple features. We first recast the two-step process as an augmented state vector approach, thereby making the error model assumptions established in chapter 5 explicit for use with traditional ensemble-based DA. In particular this clarified the requirement that the two-step method be accompanied by an alignment scheme: in short, *traditional ensemble-based DA can produce analysis positions for identified features within an estimate, but they are not able to reposition the features at those analysis positions.* We have examined this process of basis estimation, and have argued that the information contained within an ensemble is sufficient for closing a separate DA problem on feature locations. However, to access the information required (i.e., to identify the features) and to alter the ensemble members based on the DA problem (i.e., to move the features to the desired places) both require the use of methods traditionally at the fringe of the atmospheric sciences, namely image processing techniques.

Image processing and associated forms of pattern recognition are not our areas of expertise, but with minimal exploration into the field we were able to construct our own adequate alignment scheme. Together, the augmented state vector phrasing of the two-step method and the piecewise linear mapping alignment scheme have allowed us to extend the conceptual framework established in chapter 5 to more realistic problems here. In particular, we examined two simulated filtering experiments within a barotropic model. The first concerned a comparison between a strictly Lagrangian system, a point vortex system, and its Eulerian complement. Comparing our two-step method to a traditional ensemble Kalman filter, we found that the two-step method could extend the range of observational configurations under which successful filtering could occur. In some cases, the two-step method's ensemble approached the accuracy of the Lagrangian point vortex model's ensemble, the established best case scenario. The second filtering experiment we presented concerned the evolution of a barotropically unstable jet. There we showed, again, that the two-step method's ensemble remains more accurate and its individual ensemble members remain more physically plausible than when the ensemble is updated by the EnKF. Though the barotropic model in which these experiments were run is relatively simple compared to modern day numerical prediction models, we submit that the major conceptual difficulties of implementing the two-step method developed in the previous chapter have been addressed. We are hopeful that these ideas may someday find their way into an operational setting.



# Chapter 7

## Conclusions & Future Work

Here we summarize the main lines of argument and results of this thesis. We first summarize the individual chapters in order, and then we give an integrated abbreviated summary afterwards. We finish with a discussion of future work that could be continued from this thesis.

### 7.1 Chapter summaries

**Chapter 2:** This served to orient the reader so we could reinterpret state estimation and prediction informed by knowledge gained from studies in nonlinear dynamics, á la Bjerknæs's idea from a century ago brought up to date on the state of the science. We presented a directed introduction and review of the relevant basics of nonlinear dynamics so that the reader could appreciate the nonlinear dynamical interpretations of state estimation and prediction. In particular, the evolution of a specific state of a deterministic system can be traced as a trajectory through its state space: prediction concerns forecasting that trajectory, and state estimation concerns specifying the initial state. In a chaotic system monitored by noisy observations, one can never exactly specify the initial state, and therefore the entire practice is best posed probabilistically. We reviewed the theoretically correct approach to probabilistic state estimation and prediction and its various approximations that have been suggested in the literature. We ultimately ended up with a discussion on the practical (i.e., implementable) approaches to state estimation and prediction. We then reviewed the current state-of-the-art in operational numerical weather prediction efforts, and concluded that ensemble-based Kalman filters seem the most promising probabilistic state estimation implementations to date. The community has already accepted

the importance of probabilistic forecasting, as evidenced by the applications of various operational centers' Ensemble Prediction Systems. Hence, operational ensemble forecasts are already available. This especially makes the prospect of folding the information contained within those ensembles back into the state estimation process a realistic one. Ensemble-based state estimation, like the ensemble Kalman filter, is state-of-the-art, and is likely to be implemented operationally very soon.

**Chapter 3:** Here we sought to explore the consequences of generic sources of nonlinear error growth (and hence non-Gaussian PDFs) as filtered by the state-of-the-art state estimation methods. In particular, we sought to understand the differences between two classes of ensemble-based Kalman filters, deterministic filters and stochastic filters. We developed a geometric understanding of each filter's update within systems whose state spaces are small enough to visualize, and thus we were able to characterize the manner in which each filter eventually fails as nonlinear error growth is allowed to become more significant. We evaluated several ensemble assessment measures through a hierarchy of models and found for each filter that the path to filter divergence is marked by a characteristic ensemble assessment signature, thus establishing consistency in higher dimensioned models with our geometric understanding gained by examining simple models. This also served to validate the approach set forth in chapter 2 of considering the geometry of a system's state space for state estimation purposes.

**Chapter 4:** He we shifted the focus from chapter 3's investigation of generic sources of nonlinear error growth to a geophysically relevant source, namely the mis-positioning of coherent features. We first argued why errors from the mis-positioning of coherent features should be relevant to numerical weather prediction applications, and then we reviewed historical efforts to address feature position errors. This review included a brief overview of tropical cyclone prediction since it stands as a good case study of a field that has had to learn to cope with the deleterious effects of position errors. We then presented experiments comparing the kinematic updates of a system whose state vector is comprised of position information, namely a point vortex model, versus the updates of the gridpoint representation of those state vectors (i.e., we compare the Lagrangian and Eulerian updates of the same state vector information). We found that traditional ensemble-based Kalman filters have no trouble updating position information when the filters are applied directly to the positions; however, ensemble-based Kalman filters were observed to have much trouble correcting feature positions when

the positions must be inferred from a fluid field. Also, the point vortex model has given insight to when coherent features should be a problem for state estimation, namely when position errors are comparable the features' length scales. This point will be considered more carefully in the next chapter. The closing section considered a hypothetical approach to DA based on a geometrically-minded numerical method called "the level set method." Before specifying all the details of such a scheme, we were able to see that it would not generalize to more realistic problems. However, when applied to a relatively simple problem that satisfied the condition that its position errors be comparable to its feature's length scale, it was shown that improvements could be made over the standard ensemble-based approaches. Plus, it provided a useful visualization of a possible geometric approach to state estimation outside of the system's state space.

**Chapter 5:** In this chapter, the concept of alternative error models was suggested as a means to redefine estimation problems with non-Gaussian errors so that familiar, near-optimal methods may still be successfully applied. The specific example of a mixed error model including both alignment errors and additive errors was examined. Using the specific form of a soliton, an analytical solution to the KdV equation, the total errors of states following the mixed error model were shown to be non-Gaussian, and an ensemble of such states was shown to be handled poorly by a traditional EnKF, even if position observations were included. We further showed that the tangent linear hypothesis can easily be violated by incurred feature position errors. By exploring the state estimation of ensemble states realized from a mixed error model including both additive and position errors, we arrived naturally at a two-step approach to state estimation *when adhering to the implementable linear/Gaussian state estimation methods in use today*. Again taking an ensemble approach for the soliton states, this two-step approach showed a great improvement over the traditional EnKF approach. Not only did it obtain a more accurate forecast (by what ever reasonable accuracy measure one chooses to invoke), but it was able to maintain a physically plausible ensemble, each member of which is roughly statistically indistinguishable from truth. For the alignment error correction step, a very simple EnKF scheme has been used that considers only position observations. Lastly, we considered an approximation to the two-step approach that served to shift the "badness" of the analysis: instead of having un-physical updates within the individual ensemble members, one could accept physically sound individual members with an ensemble mean no longer well-approximating the minimum variance estimate.

**Chapter 6:** In this chapter we extended the simple, proof-of-concept examples from chapter 5 to spatially-extended, two-dimensional systems with multiple features. This extension was non-trivial as one is forced to invoke use of what we term an alignment scheme. We began the chapter by remolding the alternative error framework from chapter 5 into an augmented state vector approach, where position information is explicitly appended to the model state vectors. As long as the error PDFs in position were roughly Gaussian, this allowed for the simultaneous state estimation of multiple feature positions by traditional, implementable ensemble-based methods. However, the obtained analysis increments within the augmented state vector approach were not internally consistent, meaning that the updated state vectors did not have their analysis positions at the locations specified by the updated appended position information. Hence, we found we needed an alignment scheme to transform the prior state estimates so that their features' positions were at the locations specified by the first step of the approach. Once the initial states were realigned in this manner, then much of the non-Gaussianity due to position errors was removed, thus allowing successful application of the second step, traditional additive error ensemble-based state estimation. We discussed the notions of basis estimation (essentially updating position information) and the alignment scheme. We also briefly addressed common implementation concerns like how features can be identified. We presented results from two different simulated filtering exercises, both using a barotropic model. The first model revisited the kinematic updates of the point vortex model examined in chapter 4, and included the missing dynamic steps so that the state estimation information could be cycled. We found that our proposed approach and alignment scheme showed great improvement over the traditional approach. Then, in a less contrived setting than point vortices, we attempted to filter an unstable jet profile breaking up into vorticity patches. Again, our methods performed well and showed a marked improvement over the traditional approach.

## 7.2 Overall summary

This thesis has considered the effects of nonlinear error growth and its concomitant non-Gaussian uncertainty structures on ensemble-based state estimation methods. Ensemble-based extensions to the Kalman filter are shown to give accurate forecast distributions with reliable probabilistic interpretations when their assumptions of nearly linear error growth dynamics and nearly Gaussian uncertainty structures are met. However, most applications

of such extensions are inevitably for nonlinear dynamical systems, the atmosphere being one of them. The use of an ensemble allows ensemble-based filtering to extend into a regime of nonlinear error dynamics, but only weakly so. This thesis has argued that the nearly ubiquitous presence of coherent features within geophysical fluids constitutes a likely source of significant nonlinear error growth, particularly as the community refines its predictions to finer scales (e.g., the mesoscale). To address this, a potentially implementable two-step approach to state estimation was developed based on the notion that the position errors of coherent features are often Gaussian in a Lagrangian sense though they are demonstrably non-Gaussian in an Eulerian sense. This alternative error model approach was made explicit by an augmented state vector phrasing of an ensemble-based Kalman filter. It was found that a linear/Gaussian state estimation scheme was able to give meaningful analysis positions for embedded coherent features, but it was not able to reposition the features to those places. Hence, it was determined an alignment scheme was required to accomplish the repositioning. Having repositioned the features to their analysis positions, the original Eulerian non-Gaussian errors were often reduced to nearly Gaussian errors, thus allowing for the application of a standard linear/Gaussian state estimation scheme. Philosophically, this method is trying to capitalize on information that can be obtained about the state in its physical space (e.g., geometric and morphological information about the physical field) so as to compensate for information that simplified state estimation schemes discard in the system's state space. This thesis has demonstrated that such an approach is possible, and moreover, effective. It relies upon an ensemble approach to access traditionally inaccessible information about the uncertainty of a state estimate's shape. The developed approach stands as a physically-informed approach to addressing a particular form of nonlinear error growth — ground has been gained in state estimation by consideration of the phenomena at hand.

### 7.3 Continuing directions

The work within this thesis suggests numerous directions one might head with the knowledge gained. Here we suggest other fields where the ideas of alternative error models could be applied. We also address a potential issue that was raised in the final stages of our work, namely the danger of reusing observations in the second step of our method that were also used in the first step. We explain the potential issues, and outline a possible solution. This could have an immediate impact on tropical cyclone prediction, and proposals have already been written along these lines.

### 7.3.1 Logical extensions

Perhaps the most obvious direction to head next is operational implementation of the two-step approach. Indeed, scientists at the Naval Research Laboratory (NRL) in Monterey, CA, have expressed interest in this idea. Conceptually, there are two foreseeable advancements that must be made: 1) geometric descriptors beyond point locations need to be developed so that curvilinear features (e.g., fronts) and the vertical structure of features can be identified and diagnosed from ensemble forecasts, and 2) a complementary, operational-grade alignment scheme must be developed. We note that the two-step approach will most likely need not be used at every assimilation time. With an ensemble of states, one can evaluate the error linearity condition developed in chapter 5 to diagnose whether the two-step method is likely to be necessary.

In a potentially simpler venture, we have envisioned this philosophy being straightforwardly implemented in a nested-grid model setting, where one can align (essentially shift) the whole inner-grid (attending to the boundaries) as opposed to having an alignment scheme change the relative spacing of the inner-grid's gridpoints.

### 7.3.2 Possible applications

Generically, this study stands as a successful example of state vector augmentation. The ensemble context provided a natural means to forge statistically meaningful relationships between quantities otherwise difficult to link. The success suggests attempting ensemble augmentation elsewhere. We have in mind other notoriously nonlinearly related quantities like model precipitation, model parameters, and Ertel's potential vorticity.

Dr. Craig Bishop of NRL has suggested that this machinery (and overall philosophy) should also work in the reverse direction, that is, not as a method to correct ensembles but as a method to generate ensembles. Dr. Bishop has interest in running *very* large ensembles ( $O(10^4)$  members) for short times ( $O(6)$  hours) to obtain a better conditioned, state-dependent approximate covariance matrix that can be used in an operational 3DVar setting. One problem is quickly generating so many ensembles, and, in principle, a method such as ours could accomplish this in a meaningful manner, introducing relevant position errors that other quick-to-obtain methods would not.

Another possible application of the alternative error model methodology is in the area of observational quality control. As demonstrated in the EKF example in figure 5-9, potentially large model-data misfits can occur in scenarios with coherent features. Instead of outrightly rejecting the offending observation, an alternative error model approach might be adopted

so that the position content of the observation could be checked.

Relatedly, issues in forecast verifications are likely to benefit from these ideas. There is evidence that feature-based notions have already infiltrated the verification community (e.g., Nachamkin 2004), but further work could surely be of use.

### 7.3.3 Ensemble compositing

One observation one could make of the outcomes of the alignment scheme examples covered in section 6.2.2 (*cf.*, figures 6-2 and 6-5) is that there is always sizable distortion of the spatial basis near the interior tie-point. If the tie-point is found from a strong, coherent feature whose representation is distributed over a neighborhood of gridpoints, then one may feel uncomfortable with this distortion, particularly in the vicinity of a feature which is presumably sensitive to its specific structure. The reason this occurs is because we have only ascribed one tie-point to the feature. To minimize the distortion, one could try to establish more tie-points in the feature's vicinity. The problem with this is that it is not immediately clear how to invent landmark data from areas of a fluid state estimate that are not able to be uniquely identified. To this, we offer a solution we term *ensemble compositing*, which is very similar to compositing techniques used throughout the literature that produce a *typical* field rather than a strict mean field. Composites inevitably require defining a reference point (the point can be in space or time, depending on the application) about which the data being composited is oriented. Having identified corresponding interior tie-points across the ensemble, a natural reference point exists. One then transforms each ensemble member to a common grid where the interior tie-point of interest is at the origin. By re-orienting each member about this reference tie-point, one can use the ensemble to statistically determine the surrounding feature's structure. This is philosophically similar to the proposed collocation approximation of the two-step method discussed in section 5.4.

Once the referencing and re-orienting is complete, then the composite is simply the mean of the re-oriented ensemble. An ensemble mean of this oriented ensemble in the neighborhood of the tie-point will be a representative typical feature structure, however, it is the ensemble covariance, or equivalently, the cross-correlation, of this oriented ensemble which we believe can justify adding tie-points in the vicinity of a feature's position. If one finds that the cross-correlation of points neighboring the center landmark point is close to 1, then one has evidence to believe that one should not alter the spatial relationship between the elements. This is an assertion that the mean feature structure does not vary enough to warrant giving points nearby the center landmark point the freedom to move relative to that point; conversely, any point whose cross-correlation with the landmark point is low enough

can justifiably be moved with respect to that landmark point. This is all to say that an ensemble can not only give meaningful information about position discrepancy, but it also allows a statistical description of the entire field's shape. This is very similar to examining the decorrelation scale from auto-correlation functions. Ideally, one would set the threshold for cross-correlation at a statistically defined significance level, though this would require some prior estimation of effective degrees of freedom (Zwiers and von Storch 1995).

These additional tie-points will not be wholly new degrees of freedom for the DA system to reposition, rather they will accompany the landmark tie-point associated with the position of the feature. Their addition serves to extend the influence of the landmark tie-point by freezing the relative orientation of gridpoints in the neighborhood of a feature. In practice, this means that the alignment scheme will not just move a single point to match the analysis position, but a block of points whose values are highly correlated. Figure 7-1 shows two hypothetical scenarios for the  $9 \times 9$  spatial basis example used in section 6.2.2. The top two panels are for the scenario when the four closest points to the interior tie-point (the classic 5-point stencil) were determined statistically significant enough to freeze with respect to the central point. The left panel shows the new Delaunay triangulation with these four additional tie-points, and the right panel shows the updated grid from using piecewise linear mapping. The gray shaded regions indicate the areas that has been frozen with respect to one another. The bottom two panels show the same plots, only a 9-point stencil is assumed, that is, all eight surrounding gridpoints are frozen with respect to the central point. We stress that though this operation may appear to be an *ad hoc* method to decrease distortion near the feature, it is based wholly on statistical information derived from the ensemble. If a feature is extremely compact compared to the grid resolution (i.e., a  $\delta$ -function), then ensemble compositing should indicate that no neighboring gridpoints should be frozen with respect to the landmark point.

### 7.3.4 Reusing observations

An issue raised during the completion of this thesis was the potential danger in reusing observations. The two-step method as presented in section 5.4 used only position observations for basis estimation and only station observations for the second step of normal amplitude state estimation. This partitioning of the observations was made out of convenience for clarity of the example — Lagrangian observations for Lagrangian estimation and Eulerian observations for Eulerian estimation. However, one need not be constrained to make this partition. Indeed, if position observations are not available, one cannot make this partition. In section 6.1, we rephrased the two-step method as an augmented state vector

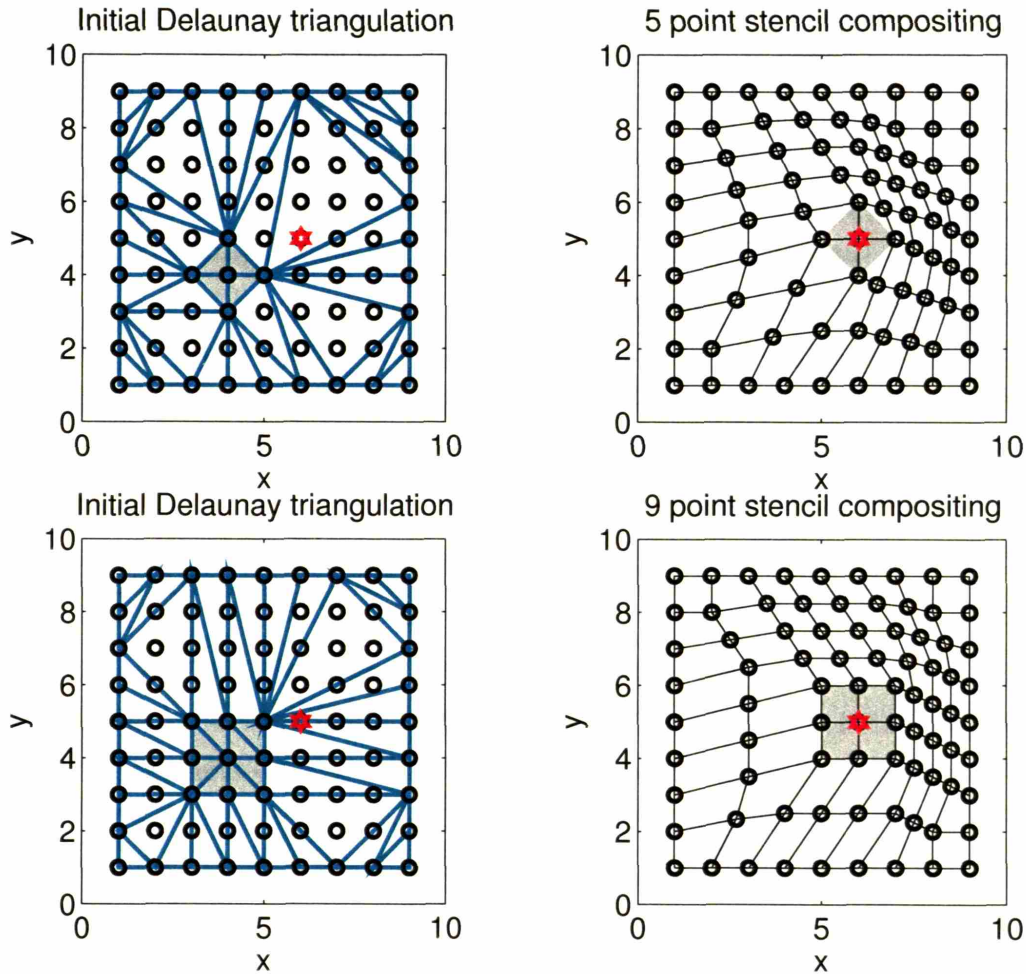


Figure 7-1: The top left shows the Delaunay triangulation for a spatial basis with five interior tie-points whose orientation has been frozen with respect to one another, denoted by the gray shaded region. The red star shows the analysis position location for the center tie-point. The right panel shows the resulting aligned grid from piecewise linear mapping. The bottom left shows the same as the top left, only for a spatial basis with nine interior tie-points. The bottom right shows the resulting aligned grid.

approach in part because it allowed a straightforward way to avoid making the partition in observations used in section 5.4. By explicitly approximating the cross-covariance between the state estimates,  $\mathbf{z}_j$ , and the position information,  $\mathbf{p}_j$ , with an augmented ensemble, we achieved the linear regression necessary to allow innovations in one component affect the other. All available observations (position and station observations) were used in each step of the two-step method.

What matters most in the Kalman filter (and other equivalent methods) is the covariance between the estimate at the estimate gridpoints and the observation locations (the  $\mathbf{PH}^T$

term in the Kalman gain matrix, see equation (2.39)), and though the estimate gridpoint locations are being changed by the alignment scheme, the observation locations are not. If the prior estimate at the observation locations changes at all as a result of the first step, then some of the information content of the observations has already been used for the estimation problem. To reuse it in the second step would be to draw the estimate closer to the observations than their uncertainty merits. One should properly account for the information already used in the observations.

This bias toward the observations can be seen by considering the derivation of the Kalman filter we presented in section 2.3.3. Equation (2.26) gives the expression for the Kalman gain matrix when allowing for correlations between the expected errors in the observations and the prior estimate. We repeat it here for convenience:

$$\mathbf{K} = (\mathbf{P}^f \mathbf{H}^T - \mathbf{G}) (\mathbf{H} \mathbf{P}^f \mathbf{H}^T + \mathbf{R} - \mathbf{H} \mathbf{G} - \mathbf{G}^T \mathbf{H}^T)^{-1}, \quad (7.1)$$

where  $\mathbf{P}^f$  is the error covariance matrix of the prior estimate's PDF,  $\mathbf{H}$  is the linearized observation operator,  $\mathbf{R}$  is the error covariance matrix of the observations' PDF, and  $\mathbf{G}$  is the expected covariance between errors in the prior estimate and errors in the observations. Strictly,  $\mathbf{G} = \langle \boldsymbol{\epsilon}^f \boldsymbol{\epsilon}^o{}^T \rangle$ , where the angle brackets denote the expectation operator and the  $\boldsymbol{\epsilon}$ 's are the respective error vectors of the prior and the observations. In most physical problems, it is unlikely that errors in the prior are correlated to errors in the observations, so  $\mathbf{G}$  is often zero. However, as discussed in section 2.3.3, if one tries to further update an analysis by the same observations used to produce it, then inclusion of the  $\mathbf{G}$  term correctly predicts a zero analysis increment. If one does not include the  $\mathbf{G}$  term, then the further analysis is drawn closer to the observations, and it has a covariance matrix indicating a more accurate analysis than the standard expression for  $\mathbf{P}^a$  with  $\mathbf{G} = \mathbf{0}$  (see equation (2.38)).

In the case of the two-step method, it is not as obvious that these issues arise because of the use of the alignment scheme. To test its importance we intend to explore the feasibility of estimating the  $\mathbf{G}$  term directly from the ensemble. In the stochastic posing of the ensemble Kalman filter (see section 3.2), one needs to make explicit perturbations of the assumed observational error covariance for each ensemble member in order to achieve the desired error statistics. Though some have argued that stochastic filters are introducing an unnecessary source of sampling error by doing this (e.g., Whitaker and Hamill 2002), we believe that those explicit perturbations could provide a way to estimate  $\mathbf{G}$ . Very preliminary experiments indicate that the impact of the  $\mathbf{G}$  correction would be small in the barotropic OSSEs shown in chapter 6, so our qualitative results are robust. We are anxious to explore this issue in further work.

# Appendix A

## Skewness and kurtosis

### A.1 Definitions

There are several common definitions for skewness and kurtosis, but all entail normalizing some measure of a distribution's higher moments by an appropriate power of the distribution's standard deviation. Here we define skewness as the unbiased estimate of a distribution's third central moment divided by the cube of its standard deviation and kurtosis as the unbiased estimate of a distribution's fourth central moment divided by the fourth power of its standard deviation less three. Consider a univariate ensemble  $z_j$  where  $j = 1, N$ . The mean, variance, and the third and fourth central moments of this ensemble are:

$$\begin{aligned}\bar{z} &= \frac{1}{N} \sum_{j=1}^N z_j & \sigma^2 &= \frac{1}{N-1} \sum_{j=1}^N (z_j - \bar{z})^2 \\ \mu_3 &= \frac{1}{N-1} \sum_{j=1}^N (z_j - \bar{z})^3 & \mu_4 &= \frac{1}{N-1} \sum_{j=1}^N (z_j - \bar{z})^4.\end{aligned}$$

These give the following expressions for skewness and kurtosis:

$$\begin{aligned}\text{skewness} &= \frac{\mu_3}{\sigma^3} \\ \text{kurtosis} &= \frac{\mu_4}{\sigma^4} - 3.\end{aligned}$$

As a Gaussian distribution is symmetric, all odd central moments are zero, and hence the expected skewness for a Gaussian is zero. All even central moments for a Gaussian are completely determined by the value of its variance. A univariate Gaussian's fourth central moment is  $3\sigma^4$ . Hence, our definition of kurtosis subtracts three so that a Gaussian has

zero kurtosis.

## A.2 Maximum Values

As skewness and kurtosis are non-dimensional quantities, it turns out that there are maximum values they can attain for a finite ensemble size  $N$ . The distribution which leads to these maximum values may be devised intuitively. Proving this, however, is a nonlinear optimization problem most easily solved by numerical methods. We offer the following as the distribution with maximal skewness and kurtosis:

$$z_j = \begin{cases} z_0 & \text{if } 1 \leq j \leq N - 1 \\ z_1 & \text{if } j = N. \end{cases}$$

This distribution gives the following central moments:

$$\begin{aligned} \bar{z} &= \frac{(N-1)z_0 + z_1}{N} & \sigma^2 &= \frac{1}{N}(z_1 - z_0)^2 \\ \mu_3 &= \frac{N-2}{N^2}(z_1 - z_0)^3 & \mu_4 &= \frac{N^2 - 3N + 3}{N^3}(z_1 - z_0)^4. \end{aligned}$$

Using these with our above definitions of skewness and kurtosis produces the following expressions:

$$\begin{aligned} \text{skewness} &= \frac{N-2}{\sqrt{N}} \\ \text{kurtosis} &= \frac{N^2 - 6N + 3}{N}. \end{aligned}$$

Note that these values do not depend on the values of  $z_0$  and  $z_1$ , only on the ensemble size. Our claim is that these are the maximum values of skewness and kurtosis that a univariate distribution of  $N$  members can attain (as skewness includes an odd moment, it can just as easily be negative, in which case the above is also the minimum possible value for skewness). This has been confirmed through exhaustive searching for relatively small distribution sizes (5 members for skewness only and 6 members for both measures). Though this does not constitute a proof, the retrieved cost functions appear rather smooth, and we suspect it is unlikely that adding further degrees of freedom will create significant local minima. Also, it should be noted that the ensembles considered above do often assume such configurations, and that these configurations have the highest skewness and kurtosis values we find.

# Appendix B

## Expected total error statistics

To evaluate the expected total error statistics shown in equations (5.9) and (5.10), and to evaluate the third central moment,  $\mathbf{\Gamma}^f$ , use the expectations of the individual errors and the Taylor series relation in equation (5.8). If  $\boldsymbol{\varepsilon}_A$  is a Gaussian random vector from the distribution  $N(\mathbf{0}, \mathbf{P}_A)$  and  $\varepsilon_D$  a Gaussian random variable from distribution  $N(0, \sigma_D^2)$ , then for the mean:

$$\begin{aligned}
 \boldsymbol{\mu}^f &= \mathbf{E}(\boldsymbol{\varepsilon}^f) = \mathbf{E}\left(\varepsilon_D \frac{d\mathbf{x}^f}{ds} + \varepsilon_D^2 \frac{1}{2} \frac{d^2\mathbf{x}^f}{ds^2} + \varepsilon_D^3 \frac{1}{3!} \frac{d^3\mathbf{x}^f}{ds^3} + \dots + \boldsymbol{\varepsilon}_A\right) \\
 &= \mathbf{E}(\varepsilon_D) \frac{d\mathbf{x}^f}{ds} + \mathbf{E}(\varepsilon_D^2) \frac{1}{2} \frac{d^2\mathbf{x}^f}{ds^2} + \mathbf{E}(\varepsilon_D^3) \frac{1}{3!} \frac{d^3\mathbf{x}^f}{ds^3} + \dots + \mathbf{E}(\boldsymbol{\varepsilon}_A) \\
 &= 0 \cdot \frac{d\mathbf{x}^f}{ds} + \sigma_D^2 \cdot \frac{1}{2} \frac{d^2\mathbf{x}^f}{ds^2} + 0 \cdot \frac{1}{3!} \frac{d^3\mathbf{x}^f}{ds^3} + \dots + \mathbf{0} \\
 &= \frac{\sigma_D^2}{2} \frac{d^2\mathbf{x}^f}{ds^2} + \frac{\sigma_D^4}{8} \frac{d^4\mathbf{x}^f}{ds^4} + \dots
 \end{aligned}$$

Similarly, for the covariance:

$$\begin{aligned}
 \mathbf{P}^f &= \mathbf{E}(\boldsymbol{\varepsilon}^f \boldsymbol{\varepsilon}^{f\text{T}}) \\
 &= \mathbf{E}\left(\left(\varepsilon_D \frac{d\mathbf{x}^f}{ds} + \varepsilon_D^2 \frac{1}{2} \frac{d^2\mathbf{x}^f}{ds^2} + \varepsilon_D^3 \frac{1}{3!} \frac{d^3\mathbf{x}^f}{ds^3} + \dots + \boldsymbol{\varepsilon}_A\right) \left(\varepsilon_D \frac{d\mathbf{x}^f}{ds} + \dots + \boldsymbol{\varepsilon}_A\right)^{\text{T}}\right) \\
 &= \mathbf{E}(\varepsilon_D^2) \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^{f\text{T}}}{ds} + \mathbf{E}(\varepsilon_D^3) \frac{1}{2} \left(\frac{d\mathbf{x}^f}{ds} \frac{d^2\mathbf{x}^{f\text{T}}}{ds^2} + \frac{d^2\mathbf{x}^f}{ds^2} \frac{d\mathbf{x}^{f\text{T}}}{ds}\right) + \mathbf{E}(\varepsilon_D^4) \frac{1}{3!} \left(\frac{d\mathbf{x}^f}{ds} \frac{d^3\mathbf{x}^{f\text{T}}}{ds^3}\right. \\
 &\quad \left. + \frac{3}{2} \frac{d^2\mathbf{x}^f}{ds^2} \frac{d^2\mathbf{x}^{f\text{T}}}{ds^2} + \frac{d^3\mathbf{x}^f}{ds^3} \frac{d\mathbf{x}^{f\text{T}}}{ds}\right) + \dots + \mathbf{E}(\boldsymbol{\varepsilon}_A \boldsymbol{\varepsilon}_A^{\text{T}}) + \text{cross-terms} \\
 &= \sigma_D^2 \frac{d\mathbf{x}^f}{ds} \frac{d\mathbf{x}^{f\text{T}}}{ds} + \frac{\sigma_D^4}{2} \left(\frac{d\mathbf{x}^f}{ds} \frac{d^3\mathbf{x}^{f\text{T}}}{ds^3} + \frac{3}{2} \frac{d^2\mathbf{x}^f}{ds^2} \frac{d^2\mathbf{x}^{f\text{T}}}{ds^2} + \frac{d^3\mathbf{x}^f}{ds^3} \frac{d\mathbf{x}^{f\text{T}}}{ds}\right) + \dots + \mathbf{P}_A,
 \end{aligned}$$

assuming that alignment errors and additive errors are uncorrelated (i.e., cross-terms = 0).

The third central moment is a third-order tensor, making notation slightly complicated. Hence, the expressions given here allow evaluation of a specifically-indexed element,  $\Gamma^f(s_0, s_1, s_2)$ . For convenience, notation will be streamlined by dots denoting first derivatives and double dots denoting second derivatives. Similarly, subscripts denote where its derivative is evaluated, for example,  $\dot{\mathbf{x}}_1 = \frac{d\mathbf{x}^f}{ds}(s_1)$ . Following the above procedure of expanding-out the Taylor series products,

$$\begin{aligned}\Gamma^f(s_0, s_1, s_2) &= \mathbf{E}(\varepsilon_D^3) (\dot{\mathbf{x}}_0 \dot{\mathbf{x}}_1 \dot{\mathbf{x}}_2) + \mathbf{E}(\varepsilon_D^4) \frac{1}{2} (\dot{\mathbf{x}}_0 \ddot{\mathbf{x}}_1 \dot{\mathbf{x}}_2 + \dot{\mathbf{x}}_0 \dot{\mathbf{x}}_1 \ddot{\mathbf{x}}_2 + \ddot{\mathbf{x}}_0 \dot{\mathbf{x}}_1 \dot{\mathbf{x}}_2) \\ &\quad + \dots + \mathbf{E}(\varepsilon_A(s_0) \varepsilon_A(s_1) \varepsilon_A(s_2)) + \text{cross-terms} \\ &= \frac{3\sigma_D^4}{2} (\dot{\mathbf{x}}_0 \ddot{\mathbf{x}}_1 \dot{\mathbf{x}}_2 + \dot{\mathbf{x}}_0 \dot{\mathbf{x}}_1 \ddot{\mathbf{x}}_2 + \ddot{\mathbf{x}}_0 \dot{\mathbf{x}}_1 \dot{\mathbf{x}}_2) + \mathbf{O}(\sigma_D^6) + \dots,\end{aligned}$$

again assuming that cross-terms are 0.

# Appendix C

## Lorenz latitude circle model

Strictly, the model is a set of  $J$  coupled ordinary differential equations:

$$\frac{dX_j}{dt} = (X_{j+1} - X_{j-2})X_{j-1} - X_j + F_j$$

where  $j = 1, J$ , and  $J$  is commonly taken as 40. As it is a latitude circle, the boundary conditions are cyclic (e.g.,  $X_0 = X_J$ ). These equations exhibit the common characteristics of forced-dissipative flow: nonlinear terms for advection, a linear term for dissipation, and a constant term for forcing. Accordingly, they can be made to exhibit chaotic behavior for adequate choice of  $F_j$ . The flow can be made unstable in a certain parameter range (here the only parameter is the constant forcing), and the instabilities are wave-like disturbances with westward (from high  $j$  to low  $j$ ) phase speed and westerly (from low  $j$  to high  $j$ ) group speed, thus mimicking traveling atmospheric disturbances.

The model is commonly run with  $F_j = F = 8.0$  as this value places it well within the unstable range; we use this value as well. We also use  $J = 40$  for our state dimension. Using the dissipative time scale of the model, Lorenz and Emanuel claim that a  $\Delta t$  of 0.05 is something like 6 hours in the atmosphere; we use a time step of 0.005 for a fourth order Runge Kutta integration.



# Appendix D

## Forms of numerical dissipation

### D.1 Wavenumber filter

The form of numerical dissipation used in our barotropic model when simulating point vortices (see section 6.3.2) is an exponential cut-off wavenumber filter that leaves all wavenumbers above a certain length scale untouched and transitions to complete elimination of small scales. It operates directly on the spectral element representations of quantities. This very scheme was used by Arbic and Flierl (2003).

The functional form is:

$$\text{filter} = \begin{cases} \exp\left(-\alpha(\sqrt{k^2 + l^2} - k_c)^M\right) & \text{for } \sqrt{k^2 + l^2} > k_c \\ 1.0 & \text{for } \sqrt{k^2 + l^2} \leq k_c \end{cases}, \quad (\text{D.1})$$

where  $\alpha$  and  $M$  are user-defined parameters, and  $k_c$  is a user-defined cut-off wavenumber. Guided by Arbic and Flierl (2003), we use  $\alpha = 18$ ,  $M = 8$ , and  $k_c = 0.65k_N$ , where  $k_N$  is the Nyquist wavenumber. This filter means that  $\mathcal{N} = \text{filter} \circ \hat{q}$ , where  $\circ$  denotes the Schur product (element-by-element multiplication). Figure D-1 shows the functional form in spectral space.

Using this implementation, we find that to maintain coherent Gaussian vortices, we must use a  $d$  above  $2.5\Delta x$ , where  $\Delta x$  is the gridpoint resolution (assumed uniform in  $x$  and  $y$  directions over the whole grid). The minimum  $d$  we use is  $3\Delta x$ , giving the vortices an effective diameter of  $6\Delta x$ .

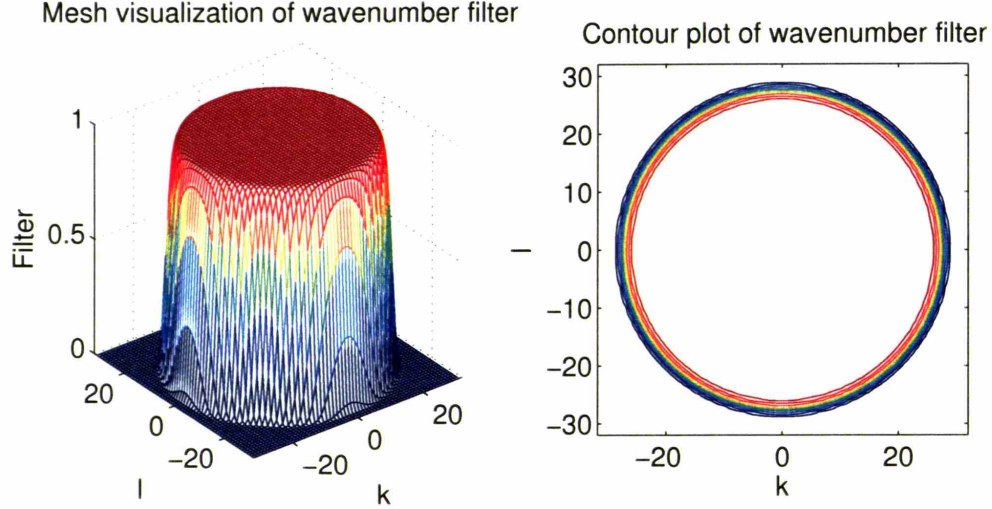


Figure D-1: The left panel shows a mesh visualization of the exponential wavenumber filter template. The  $k$  and  $l$  axes show the wavenumbers. The right panel shows the same information as a contour plot.

## D.2 Biharmonic hyperdiffusion

The form of numerical dissipation used in our barotropic model when evolving the barotropically unstable jet (see section 6.3.3) is a biharmonic hyperdiffusion term. As opposed to the wavenumber filter in the section above, this operator damps all scales. However, the fourth power makes it quite scale selective. This same term was used by Flierl et al. (1987) in their use of their barotropic model.

We implement our viscosity by a unique timestepping algorithm we learned from Flierl (personal communication, 2002). Instead of treating viscosity as a term on the RHS as in:

$$\frac{\partial q}{\partial t} = -J(\psi, q) - \nu K^4 q,$$

where  $q$  and  $\psi$  are understood to be the two-dimensional Fourier transforms and  $K = (k+l)$  is the total wavenumber, Flierl's scheme uses a multiplicative viscosity:

$$\begin{aligned} \frac{\partial q}{\partial t} + \nu K^4 q &= -J(\psi, q) \\ \frac{\partial}{\partial t} e^{\nu K^4 t} q &= -e^{\nu K^4 t} J(\psi, q). \end{aligned}$$

Now the quantity  $\exp(\nu K^4 t)q$  is what is explicitly evolved, and when the vorticity field itself is needed, one simply multiplies by  $\exp(-\nu K^4 t)$ .

# Appendix E

## Doubly periodic point vortex model

A point vortex model is a fully Lagrangian model comprised of  $2N_v$  ODEs, where  $N_v$  is the number of vortices. It is not intrinsically bounded to a domain, so to render a point vortex doubly periodic is not an obvious venture. We are aware of no recurrence relation or infinite limit that would allow this form analytically, so we were forced to approximate it numerically. We have constructed a remarkably accurate approximation that simulates double periodicity by explicitly appending “shells” of neighboring domains around the central domain of interest and by decomposing each contributing vortex into a linear superposition of three “sub-vortices.” This will be made clear below.

The dynamics of a normal point vortex are given in equations (4.4) and (4.5). They are repeated here for convenience:

$$\frac{dx}{dt}(\mathbf{x}) = - \sum_{i=1}^{N_v} \frac{\Gamma_i}{2\pi} \frac{y - y_i}{\|\mathbf{x} - \mathbf{x}_i\|^2} \quad (\text{E.1})$$

$$\frac{dy}{dt}(\mathbf{x}) = \sum_{i=1}^{N_v} \frac{\Gamma_i}{2\pi} \frac{x - x_i}{\|\mathbf{x} - \mathbf{x}_i\|^2}, \quad (\text{E.2})$$

where  $\|\cdot\|$  denotes the  $\ell_2$  norm (i.e., magnitude) of its vector argument.

The dynamics are essentially those of mutual advection, meaning every point vortex contributes to the ensuing motion of all other vortices except itself. Each vortex’s contribution to the others sums linearly.

Hence, for a non-periodic case (with an infinite domain), the contribution to the  $u$  velocity from vortex  $i$  to vortex  $j$  is

$$\frac{dx_j}{dt} = - \frac{\Gamma_i}{2\pi} \frac{y_j - y_i}{\|\mathbf{x}_j - \mathbf{x}_i\|^2} \quad (\text{E.3})$$

If one were to try to simulate double periodicity by naively adding a number of shells, say  $N_{sh}$ , around a central domain, then to amend the equation (E.3), one must consider the contributions to vortex  $j$  from all  $(2N_{sh} + 1)^2$  versions of the  $i$ th vortex. This adds a summation:

$$\frac{dx_j}{dt} = - \sum_{k=1}^{(2N_{sh}+1)^2} \frac{\Gamma_i}{2\pi} \frac{y_j - y_k}{\|\mathbf{x}_j - \mathbf{x}_k\|^2}, \quad (\text{E.4})$$

where  $k$ th positions are now related to the original  $i$  by an integer number of domain lengths.

This method will not in general approximate the truly doubly periodic version, however, because the point vortices are felt at a distance (through the elliptic operator relating vorticity to streamfunction and hence velocity), and there are imbalanced number of vortices occupying the shells (i.e., if vortex  $j$  is to vortex  $i$ 's right, then in the shell approximation, there will always be one more copy of vortex  $i$  to  $j$ 's left than its right). This imbalance, for instance, will disallow the naive shell approximation from arresting the motion of two like-signed vortices of equal strength at the positions  $(0, -L/4)$  and  $(0, L/4)$  in a square domain of length  $L$  centered at the origin — a truly periodic system would be halted.

To address this imbalance, we have found it convenient to separate the contributions from each advecting vortex into three components (“sub-vortices”). Each vortex is assigned two companions outside the original domain, one a distance  $L$  away in the  $x$  direction and one a distance  $L$  away in the  $y$  direction. The companions are located in the direction (e.g., left or right) that is farther to the domain boundary (i.e., across either the  $x$ -axis or the  $y$ -axis). The three sub-vortices are then ascribed a linear parsing of the original vortex’s circulation strength. Each of the two companion sub-vortices receives a share of the circulation in proportion to its nearness to the vortex in the central domain that is to be advected, and the sub-vortex left at the location of the original vortex is given a strength such the three sub-vortices’ strengths sum to the original strength (see equation (E.7)). Hence, once the decomposed sub-vortices are tiled in a shell pattern, they have been constructed to overlay one another and sum to their original strength in the central domain, but they are allowed to have fractional strengths outside the central domain. These unpaired fractional vortices almost exactly cancel the imbalance. Equal vortex pairs initialized at  $(0, -L/4)$  and  $(0, L/4)$  will halt under this method.

$$\Gamma_x = w_x \Gamma_i = \frac{(x_j - x_k)}{L} \Gamma_i \quad (\text{E.5})$$

$$\Gamma_y = w_y \Gamma_i = \frac{(y_j - y_k)}{L} \Gamma_i \quad (\text{E.6})$$

$$\Gamma_o = (1 - w_x - w_y) \Gamma_i \quad (\text{E.7})$$

Figure E-1 shows the agreement achieved for a lengthy integration of both the doubly periodic point vortex system and a companion barotropic model with point vortices inserted within it (that truly is doubly periodic!). As you can see, they agree very well. No adjusting to either state has been done.




Figure E-1: The circular contoured patches are overlaid successive profiles of barotropic vorticity. The colored lines are the tracks of the doubly periodic point vortex model initialized with the same vortex locations.



# Bibliography

- Aberson, S. D., 2002: Two years of operation hurricane synoptic surveillance. *Wea. Forecasting*, 17, 1101–1110.
- Alexander, G. D., Weinman, J. A., and Schols, J. L., 1998: The use of digital warping of microwave integrated water vapor imagery to improve forecasts of marine extratropical cyclones. *Mon. Wea. Rev.*, 126, 1469–1496.
- Anderson, D., Hodges, K. I., and Hoskins, B. J., 2003: Sensitivity of feature-based analysis methods of storm tracks to the form of background field removal. *Mon. Wea. Rev.*, 131, 565–573.
- Anderson, J. L., 2001: An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.*, 129, 2884–2903.
- Anderson, J. L., 2004: NCAR scientist. Personal communication.
- Anderson, J. L., 2005: A hierarchical ensemble filter for data assimilation. *Mon. Wea. Rev.*, submitted.
- Anderson, J. L. and Anderson, S. L., 1999: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Wea. Rev.*, 127, 2741–2758.
- Arbic, B. K. and Flierl, G. R., 2003: Coherent vortices and kinetic energy ribbons in asymptotic, quasi two-dimensional f-plane turbulence. *Physics of Fluids*, 15, 2177–2189.
- Aref, H., 1983: Integrable, chaotic, and turbulent vortex motion in two-dimensional flows. *Ann. Rev. Fluid Mech.*, 15, 345–389.
- Aref, H. and Pomphrey, N., 1982: Integrable and chaotic motions of four vortices I. The case of identical vortices. *Proc. Roy. Soc. London Ser. A*, 380, 359–387, issue 1779.

- Barnes, S., 1964: A technique for maximizing details in numerical weather map analysis. *J. Appl. Meteor.*, 3, 396–409.
- Bergé, P., Pomeau, Y., and Vidal, C., 1984: *Order Within Chaos*. John Wiley & Sons, Inc., 329pp.
- Berloff, P. S. and McWilliams, J. C., 2002: Material transport in oceanic gyres. Part III: Hierarchy of stochastic models. *J. Phys. Oceanogr.*, 32, 797–830.
- Bishop, C. H., Etherton, B. J., and Majumdar, S. J., 2001: Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Wea. Rev.*, 129, 420–436.
- Bjerknes, V., 1904: The problem of weather forecasting as a problem in mechanics and physics. *Meteorol. Z.*, 21, 1–7.
- Bjerknes, V. and different collaborators, 1911: *Dynamic Meteorology and Hydrography*, volume Part II. — Kinematics. Carnegie Institute of Washington, 175pp.
- Blackmon, M. L., Wallace, J. M., Lau, N.-C., and Mullen, S. L., 1977: An observational study of the Northern Hemisphere wintertime circulation. *J. Atmos. Sci.*, 34, 1040–1053.
- Blumen, W., 1975: An analytical view of updating meteorological variables: Part I. Phase errors. *J. Atmos. Sci.*, 32, 274–286.
- Brewster, K. A., 2003a: Phase-correcting data assimilation and application to storm-scale numerical weather prediction. Part I: Method description and simulation testing. *Mon. Wea. Rev.*, 131, 480–492.
- Brewster, K. A., 2003b: Phase-correcting data assimilation and application to storm-scale numerical weather prediction. Part II: Application to a severe storm outbreak. *Mon. Wea. Rev.*, 131, 493–507.
- Buizza, R., Miller, M., and Palmer, T. N., 1999: Stochastic representation of model uncertainties in the ECMWF Ensemble Prediction System. *Quart. J. Roy. Meteor. Soc.*, 125, 2887–2908.
- Burgers, G., van Leeuwen, P., and Evensen, G., 1998: Analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.*, 126, 1719–1724.

- Charney, J., Fjørtoft, R., and von Neumann, J., 1950: Numerical integration of the barotropic vorticity equation. *Tellus*, 2, 237–257.
- Charney, J., Halem, M., and Jastrow, R., 1969: Use of incomplete historical data to infer the present state of the atmosphere. *J. Atmos. Sci.*, 26, 1160–1163.
- Chin, T. M., Mariano, A. J., and Chassignet, E. P., 1999: Spatial regression and multiscale approximations for sequential data assimilation in ocean models. *J. Geophys. Res. — Oceans*, 104, 7991–8014.
- Cohn, S. E., 1993: Dynamics of short-term univariate forecast error covariances. *Mon. Wea. Rev.*, 121, 3123–3149.
- Cohn, S. E., 1997: An introduction to estimation theory. *J. Meteor. Soc. Japan*, 75, 257–288.
- D’Agostini, G., 2003: *Bayesian Reasoning in Data Analysis: A Critical Introduction*. World Scientific, 329pp.
- Dee, D. P., 1991: Simplification of the Kalman filter for meteorological data assimilation. *Quart. J. Roy. Meteor. Soc.*, 117, 365–384.
- Dee, D. P. and da Silva, A. M., 1998: Data assimilation in the presence of forecast bias. *Quart. J. Roy. Meteor. Soc.*, 124, 269–295.
- Dee, D. P., Rukhovets, L., Todling, R., da Silva, A. M., and Larson, J. W., 2001: An adaptive buddy check for observational quality control. *Quart. J. Roy. Meteor. Soc.*, 127, 2451–2471.
- Delsol, C., 2005: *Objective Identification of Weather Systems in Satellite Imagery using Fourier Shape Descriptors*. Ph.D. thesis, ESSC, University of Reading, UK, submitted.
- Delsol, C. and Hodges, K. I., 2003: Identifying mid-latitude weather systems using Fourier shape descriptors. *AGU/EGS/EGU Joint Session, Nice, France*, poster presented at conference.
- DeMaria, M., 1997: Summary of the Tropical Prediction Center/National Hurricane Center tropical cyclone track and intensity guidance models. TPC/NHC Web Site, <http://www.nhc.noaa.gov/aboutmodels.html>.
- Desroziers, G., 1997: A coordinate change for data assimilation in spherical geometry of frontal structures. *Mon. Wea. Rev.*, 125, 3030–3038.

- Desroziers, G. and Lafore, J.-P., 1993: A coordinate transformation for objective frontal analysis. *Mon. Wea. Rev.*, 121, 1531–1553.
- Dickinson, S. and Brown, R. A., 1996: A study of near-surface winds in marine cyclones using multiple satellite sensors. *J. Appl. Meteor.*, 35, 769–781.
- Drazin, P. G. and Johnson, R. S., 1989: *Solitons: an Introduction*. Cambridge, 226 pp.
- Dritschel, D. G., 1988: Contour surgery: A topological reconnection scheme for extended integrations using contour dynamics. *J. Comp. Phys.*, 77, 240–266.
- Durran, D. R., 1991: The third-order Adams-Bashforth method: An attractive alternative to leapfrog time differencing. *Mon. Wea. Rev.*, 119, 702–720.
- Ehrendorfer, M., 1997: Predicting the uncertainty of numerical weather forecasts: a review. *Meteorol. Z.*, 6, 147–183.
- Eliassen, A., 1954: Provisional report on calculation of spatial covariance and autocorrelation of the pressure field. Technical report, Inst. Weather and Climate Research, Acad. Sci., Oslo, report 5.
- Elsberry, R. L., 1995: Recent advancements in dynamical tropical cyclone track prediction. *Meteorol. Atmos. Phys.*, 56, 81–99.
- Emanuel, K. A., 1999: Thermodynamic control of hurricane intensity. *Nature*, 401, 665–669.
- Enăchescu, D. M., 1985: Monte Carlo method for nonlinear filtering. *Proceedings of the Fourth Pannonian Symposium on Mathematical Statistics (1983)*, Mathematical Statistics and its Applications, Kluwer, 125–139.
- Epstein, E., 1969: Stochastic dynamic prediction. *Tellus*, 21, 739–759.
- Evans, L. C., 1997: Partial differential equations and Monge-Kantorovich mass transfer. R. B. et al., ed., *Current Developments in Mathematics*, International Press, 26–78.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, 99, 10143–10162.
- Evensen, G. and van Leeuwen, P., 1996: Assimilation of Geosat altimeter data for the Agulhas Current using the ensemble Kalman filter with a quasi-geostrophic model. *Mon. Wea. Rev.*, 124, 85–96.

- Farrell, B. F. and Ioannou, P., 2001: State estimation using a reduced-order Kalman filter. *J. Atmos. Sci.*, 58, 3666–3680.
- Flierl, G. R., 2002: Prof. of oceanography, MIT. Personal communication.
- Flierl, G. R., Malanotte-Rizzoli, P., and Zabusky, N. J., 1987: Nonlinear waves and coherent vortex structures in barotropic  $\beta$ -plane jets. *J. Phys. Oceanogr.*, 17, 1408–1438.
- Gaspari, G. and Cohn, S. E., 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, 125, 723–757.
- Ghil, M. and Malanotte-Rizzoli, P., 1991: Data assimilation in meteorology and oceanography. *Advances in Geophysics*, volume 33, Academic Press, 141–266.
- Gilchrist, B. and Cressman, G., 1954: An experiment in objective analysis. *Tellus*, 6, 309–318.
- Goerss, J. S., 2000: Tropical cyclone track forecasts using an ensemble of dynamical models. *Mon. Wea. Rev.*, 128, 1187–1193.
- Gordon, N. J., Salmond, D. J., and Smith, A. F. M., 1993: Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F*, 140, 107–113.
- Goshtasby, A., 1986: Piecewise linear mapping functions for image registration. *Pattern Recognition*, 19, 459–466.
- Goshtasby, A., 1987: Piecewise cubic mapping functions for image registration. *Pattern Recognition*, 20, 525–533.
- Grassotti, C., Iskenderian, H., and Hoffman, R. N., 1999: Fusion of surface radar and satellite rainfall data using feature calibration and alignment. *J. Appl. Meteor.*, 38, 677–695.
- Haidvogel, D. B., Robinson, A. R., and Schulman, E. E., 1980: Accuracy, efficiency, and stability of 3 numerical-models with application to open ocean problems. *J. Comput. Phys.*, 34, 1–53.
- Haker, S., Zhu, L., Tannenbaum, A., and Angenent, S., 2004: Optimal mass transport for registration and warping. *International J. Computer Vision.*, 60, 225–240.

- Hakim, G. J., 2000: Climatology of coherent structures on the extratropical tropopause. *Mon. Wea. Rev.*, 128, 385–406.
- Hakim, G. J., 2004: Prof. of meteorology, University of Washington. Personal communication.
- Hald, O. H., 1979: Convergence of vortex methods for Euler's equations. II. *SIAM J. Numer. Anal.*, 16, 726–755.
- Haller, G. and Yuan, G., 2000: Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D*, 147, 352–370.
- Hamill, T. and Snyder, C., 2000: A hybrid ensemble Kalman filter-3D variational analysis scheme. *Mon. Wea. Rev.*, 128, 2905–2919.
- Hamill, T. M., 2001: Interpretation of rank histograms for verifying ensemble forecasts. *Mon. Wea. Rev.*, 129, 550–560.
- Hansen, J. A. and Smith, L. A., 2000: The role of operational constraints in selecting supplementary observations. *J. Atmos. Sci.*, 57, 2859–2871.
- Hansen, J. A. and Smith, L. A., 2001: Probabilistic noise reduction. *Tellus*, 53A, 587–598.
- Heming, J. T. and Radford, A. M., 1998: The performance of the United Kingdom Meteorological Office global model in predicting the tracks of Atlantic tropical cyclones in 1995. *Mon. Wea. Rev.*, 126, 1323–1331.
- Henry, A. J., 1916: 3. Weather forecasting—preliminary statement of the problem. A. J. Henry, E. H. Bowie, H. J. Cox, and H. C. Frankenfield, eds., *Weather Forecasting in the United States*, W. B. No. 583, U. S. Department of Agriculture, Weather Bureau, 69–76, 370 pp.
- Hodges, K. I., 1994: A general method for tracking analysis and its applications to meteorological data. *Mon. Wea. Rev.*, 122, 2573–2586.
- Hodges, K. I., 1995: Feature tracking on the unit sphere. *Mon. Wea. Rev.*, 123, 3458–3465.
- Hodges, K. I., 1998: Feature-point detection using distance transforms: Application to tracking tropical convective complexes. *Mon. Wea. Rev.*, 126, 785–795.
- Hodges, K. I., 1999: Adaptive constraints for feature tracking. *Mon. Wea. Rev.*, 127, 1362–1373.

- Hoffman, R. N. and Grassotti, C., 1996: A technique for assimilating SSM/I observations of marine atmospheric storms: Tests with ECMWF analyses. *J. Appl. Meteor.*, 35, 1177–1188.
- Hoffman, R. N., Liu, Z., Louis, J.-F., and Grassotti, C., 1995: Distortion representation of forecast errors. *Mon. Wea. Rev.*, 123, 2758–2770.
- Hoskins, B. J., 1975: The geostrophic momentum approximation and the semi-geostrophic equation. *J. Atmos. Sci.*, 32, 233–242.
- Hoskins, B. J. and Hodges, K. I., 2002: New perspectives on the Northern Hemisphere winter storm tracks. *J. Atmos. Sci.*, 59, 1041–1061.
- Hoskins, B. J., McIntyre, M. E., and Roberston, A. W., 1985: On the use and significance of isentropic potential vorticity maps. *Quart. J. Roy. Meteor. Soc.*, 111, 877–946.
- Houtekamer, P. L. and Mitchell, H. L., 1998: Data assimilation using an ensemble Kalman filter technique. *Mon. Wea. Rev.*, 126, 796–811.
- Houtekamer, P. L. and Mitchell, H. L., 2001: A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, 129, 123–137.
- Houtekamer, P. L., Mitchell, H. L., and Pellerin, G., 2004: Atmospheric data assimilation with an ensemble Kalman filter: Results with real observations. *Mon. Wea. Rev.*, submitted.
- Ide, K. and Ghil, M., 1997a: Extended Kalman filtering for vortex systems. Part I: Methodology and point vortices. *Dyn. Atmos. Oceans*, 27, 301–332.
- Ide, K. and Ghil, M., 1997b: Extended Kalman filtering for vortex systems. Part II: Rankine vortices and observing-system design. *Dyn. Atmos. Oceans*, 27, 333–350.
- Ide, K., Courtier, P., Ghil, M., and Lorenc, A., 1997: Unified notation for data assimilation: Operational, sequential and variational. *J. Meteor. Soc. Japan*, 75, 181–189.
- Ikeda, K., 1979: Multiple-valued stationary state and its instability of the transmitted light by a ring cavity system. *Optics Communications*, 30, 257–261.
- Jazwinski, A. H., 1970: *Stochastic Processes and Filtering Theory*. Academic Press, 376 pp.
- Judd, K. and Smith, L. A., 2001: Indistinguishable states. I. Perfect model scenario. *Physica D*, 151, 125–141.

- Kalman, R. E., 1960: A new approach to linear filtering and prediction problems. *Trans. ASME — J. Bas. Eng.*, 82, 35–45.
- Kim, S., Eyink, G. L., Restrepo, J. M., Alexander, F. J., and Johnson, G., 2003: Ensemble filtering for nonlinear dynamics. *Mon. Wea. Rev.*, 131, 2586–2594.
- Krishnamurti, T. N., Kishtawal, C. M., Zhang, Z., LaRow, T., Bachiochi, D., and Williford, E., 2000: Multimodel ensemble forecasts for weather and seasonal climate. *J. Climate*, 13, 4196–4216.
- Kurihara, Y., Tuleya, R. E., and Bender, M. A., 1998: The GFDL hurricane prediction system and its performance in the 1995 hurricane season. *Mon. Wea. Rev.*, 126, 1306–1322.
- Kuznetsov, L., Ide, K., and Jones, C. K. R. T., 2003: A method for assimilation of Lagrangian data. *Mon. Wea. Rev.*, 131, 2247–2260.
- Landsea, C. and Lawrence, M., 2004: How accurate are the forecasts from the National Hurricane Center? AOML Web Site, <http://www.aoml.noaa.gov/hrd/tcfaq/F7.html>.
- Lawson, W. G. and Hansen, J. A., 2004: Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth. *Mon. Wea. Rev.*, 132, 1966–1981.
- Lawson, W. G. and Hansen, J. A., 2005: Alignment error models and ensemble-based data assimilation. *Mon. Wea. Rev.*, in press.
- Legrand, S., Legat, V., and Deleersnijder, E., 2000: Delaunay mesh generation for an unstructured-grid ocean general circulation model. *Ocean Modelling*, 2, 17–28.
- Leith, C., 1974: Theoretical skill of Monte Carlo forecasts. *Mon. Wea. Rev.*, 102, 409–418.
- LeMarshall, J. F., 1998: Cloud and water vapour motion vectors in tropical cyclone track forecasting—a review. *Meteorol. Atmos. Phys.*, 65, 141–151.
- Leslie, L. M. and Fraedrich, K., 1990: Reduction of tropical cyclone position errors using an optimal combination of independent forecasts. *Wea. Forecasting*, 5, 158–161.
- Leslie, L. M. and Holland, G. J., 1995: On the bogussing of tropical cyclones in numerical models: A comparison of vortex profiles. *Meteorol. Atmos. Phys.*, 56, 101–110.

- Leslie, L. M., Abbey, Jr., R. F., and Holland, G. J., 1998a: Tropical cyclone track predictability. *Meteorol. Atmos. Phys.*, 65, 223–231.
- Leslie, L. M., LeMarshall, J. F., Morison, R. P., Spinoso, C., Purser, R. J., Pescod, N., and Seecamp, R., 1998b: Improved hurricane track forecasting from the continuous assimilation of high quality satellite wind data. *Mon. Wea. Rev.*, 126, 1248–1257.
- Li, Y. and Sattinger, D. H., 1998: Matlab codes for nonlinear dispersive wave equations. Technical report, Univ. of Minnesota, available at <http://www.math.usu.edu/~dhs/codes.ps>.
- Li, Y., Ménard, R., Riishøjgaard, L. P., Cohn, S. E., and Rood, R. B., 1998: A study on assimilating potential vorticity data. *Tellus*, 50A, 490–506.
- Liebelt, P. B., 1967: *An Introduction to Optimal Estimation*. Addison-Wesley, 273pp.
- Liu, Y., Zhang, D., and Yau, M. K., 1997: A multiscale numerical study of Hurricane Andrew (1992). Part I: Explicit simulation and verification. *Mon. Wea. Rev.*, 125, 3073–3093.
- Lorenç, A., 1986: Analysis methods for numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, 112, 1177–1194.
- Lorenz, E. N., 1963: Deterministic nonperiodic flow. *J. Atmos. Sci.*, 20, 130–141.
- Lorenz, E. N., 1965: A study of the predictability of a 28-variable atmospheric model. *Tellus*, 17, 321–333.
- Lorenz, E. N., 1969: The predictability of a flow which possesses many scales of motion. *Tellus*, 21, 289–307.
- Lorenz, E. N., 1995: Predictability – A problem partly solved. *Predictability*, ECMWF, Seminar Proceedings, Shinfield Park, Reading, RG2 9AX, 1–18.
- Lorenz, E. N., 2004: Emeritus prof. of meteorology, MIT. Personal communication.
- Lorenz, E. N. and Emanuel, K. A., 1998: Optimal sites for supplementary weather observations: Simulation with a small model. *J. Atmos. Sci.*, 55, 399–414.
- Lorenz, E. N. and Krishnamurthy, V., 1987: On the nonexistence of a slow manifold. *J. Atmos. Sci.*, 44, 2940–2950.

- Luckman, A., Murray, T., Jiskoot, H., Pritchard, H., and Strozzi, T., 2003: ERS SAR feature-tracking measurement of outlet glacier velocities on a regional scale in East Greenland. *Annals of Glaciology*, 36, 129–134.
- Lyons, W. A., 1997: *The Handy Weather Answer Book*. Visible Ink Press, 397pp.
- Mackenzie, D., 2003: Ensemble Kalman filters bring weather models up to date. *SIAM News*, 36.
- Mardia, K. V., 1970: Measures of multivariate skewness and kurtosis with applications. *Biometrika*, 57, 519–530.
- Mariano, A. J., 1990: Contour analysis: A new approach for melding geophysical fields. *J. Atmos. Oceanic Tech.*, 7, 285–295.
- Mariano, A. J. and Chin, T. M., 1996: Feature and contour based data analysis and assimilation in physical oceanography. R. J. Adler, P. Müller, and B. L. Rozovskii, eds., *Stochastic Modelling in Physical Oceanography*, Birkhäuser, 311–342.
- Maybeck, P. S., 1979: *Stochastic Models, Estimation, and Control*, volume 141-1 of *Mathematics in Science and Engineering*. Academic Press, 423pp.
- McMurdie, L. and Mass, C., 2004: Major numerical forecast failures over the Northwest Pacific. *Wea. Forecasting*, 19, 338–356.
- McWilliams, J. C., 1984: The emergence of isolated coherent vortices in turbulent flow. *J. Fluid Mech.*, 146, 21–43.
- McWilliams, J. C., 1990: The vortices of two-dimensional turbulence. *J. Fluid Mech.*, 219, 361–385.
- Miller, R. N., Robinson, A. R., and Haidvogel, D. B., 1983: A baroclinic quasigeostrophic open ocean model. *J. Comput. Phys.*, 50, 38–70.
- Miller, R. N., Ghil, M., and Gauthiez, F., 1994: Advanced data assimilation in strongly nonlinear dynamical systems. *J. Atmos. Sci.*, 51, 1037–1056.
- Moore, A. M., Perez, C. L., and Zavala-Garay, J., 2002: A non-normal view of the wind-driven circulation. *J. Phys. Oceanogr.*, 32, 2681–2705.
- Morss, R. E., 1998: *Adaptive Observations: Idealized Sampling Strategies for Improving Numerical Weather Prediction*. Ph.D. thesis, Massachusetts Institute of Technology.

- Muccino, J. C. and Bennett, A. F., 2002: Generalized inversion of the Korteweg-de Vries equation. *Dyn. Atmos. Oceans*, 35, 227–263.
- Nachamkin, J. E., 2004: Mesoscale verification using meteorological composites. *Mon. Wea. Rev.*, 132, 941–955.
- Nehrkorn, T., Hoffman, R. N., Grassotti, C., and Louis, J.-F., 2003: Feature calibration and alignment to represent model forecast errors: Empirical regularization. *Quart. J. Roy. Meteor. Soc.*, 129, 195–218.
- Nolan, D. S., 2004: Vortex sheets, vortex rings, and a mesocyclone. *22nd Conference on Severe Local Storms*, American Meteorological Society, Hyannis, MA.
- Novikov, E. A. and Sedov, Y. B., 1978: Stochastic properties of a four-vortex system. *Sov. Phys. JETP*, 48, 440–444.
- Osher, S. and Sethian, J. A., 1988: Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *J. Comp. Phys.*, 79, 12–49.
- Palmer, T. N., 2003: ECMWF scientist, MIT Houghton Lecturer. Personal communication.
- Panofsky, H., 1949: Objective weather-map analysis. *J. Meteor.*, 6, 386–392.
- Parrish, D. and Derber, J., 1992: The National Meteorological Center’s spectral statistical-interpolation analysis system. *Mon. Wea. Rev.*, 120, 1747–1763.
- Penland, C., 2003: A stochastic approach to nonlinear dynamics. *Bull. Amer. Meteor. Soc.*, *on-line extended version*, 84, ES43–ES52, doi: 10.1175/BAMS-84-7-Penland.
- Pham, D., Verron, J., and Roubaud, M., 1998: A singular evolutive extended Kalman filter for data assimilation in oceanography. *J. Marine Sci.*, 16, 323–340.
- Polvani, L. and Plumb, R., 1992: Rossby wave breaking, microbreaking, filamentation, and secondary vortex formation: The dynamics of a perturbed vortex. *J. Atmos. Sci.*, 49, 462–476.
- Provenzale, A., 1999: Transport by coherent barotropic vortices. *Annu. Rev. Fluid Mech.*, 31, 55–93.
- Provenzale, A. and Balmforth, N. J., 1998: Chaos and structures in geophysics and astrophysics, woods Hole GFD Summer School notes.

- Puri, K., Barkmeijer, J., and Palmer, T. N., 2001: Ensemble prediction of tropical cyclones using targeted diabatic singular vectors. *Quart. J. Roy. Meteor. Soc.*, 127, 709–731.
- Rabier, E., Jarvinen, H., Klinker, E., Mahfouf, J., and Simmons, A., 2000: The ECMWF operational implementation of four-dimensional variational assimilation. I: Experimental results with simplified physics. *Quart. J. Roy. Meteor. Soc.*, 126, 1143–1170.
- Ravela, S., Emanuel, K. A., and McLaughlin, D., 2004: Data assimilation by field alignment. Submitted.
- Reichle, R. H., McLaughlin, D. B., and Entekhabi, D., 2002: Hydrologic data assimilation with the ensemble Kalman filter. *Mon. Wea. Rev.*, 130, 103–114.
- Sethian, J. A., 1996: *Level Set Methods and Fast Marching Methods*. Cambridge University Press, 378pp.
- Silverman, B. W., 1986: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 175pp.
- Smith, L. A., 2002: What might we learn from climate forecasts? *Proceedings of the National Academy of Sciences*, 99, suppl. 1, 2487–2492.
- Smith, L. A. and Hansen, J. A., 2003: Extending the limits of forecast verification with the MST. *Mon. Wea. Rev.*, submitted.
- Snyder, C., 2004: NCAR scientist. Personal communication.
- Strogatz, S. H., 1994: *Nonlinear Dynamics and Chaos*. Addison-Wesley Publishing Company, 498pp.
- Surgi, N., Pan, H.-L., and Lord, S. J., 1998: Improvement of the NCEP global model over the tropics: An evaluation of model performance during the 1995 hurricane season. *Mon. Wea. Rev.*, 126, 1287–1305.
- Tarantola, A., 1987: *Inverse Problem Theory*. Elsevier, 613pp.
- Thomson, D. J., 1987: Criteria for the selection of stochastic models of particle trajectories in turbulent flows. *J. Fluid Mech.*, 180, 529–556.
- Tippett, M. K., Anderson, J. L., Bishop, C. H., Hamill, T. M., and Whitaker, J. S., 2003: Ensemble square-root filters. *Mon. Wea. Rev.*, in Press.

- Tiwari, R. K. and Rao, K. N. N., 1999: Phase space structure, attractor dimension, lyapunov exponent and nonlinear prediction from earth's atmospheric angular momentum time series. *Pure and Applied Geophysics*, 156, 719–736.
- Togneri, R. and Deng, L., 2003: Joint state and parameter estimation for a target-directed nonlinear dynamic system model. *IEEE Trans. Sig. Processing*, 51, 3061–3070.
- Tsonis, A. A. and Elsner, J. B., 1989: Chaos, strange attractors, and weather. *Bull. Amer. Meteor. Soc.*, 70, 14–23.
- van Leeuwen, P. J., 2003: A variance-minimizing filter for large-scale applications. *Mon. Wea. Rev.*, 131, 2071–2084.
- Velden, C. S., Hayden, C. M., Nieman, S. J., Menzel, W. P., Wanzong, S., and Goerss, J. S., 1997: Upper-tropospheric winds derived from geostationary satellite water vapor observations. *Bull. Amer. Meteor. Soc.*, 78, 173–195.
- Verlaan, M. and Heemink, A. W., 1997: Tidal flow forecasting using reduced rank square root filters. *Stochastic Hydrology and Hydraulics*, 11, 349–368.
- Watson, D. F., 1992: *Contouring: a Guide to the Analysis and Display of Spatial Data*. Pergamon Press, 321pp.
- Weber, H. C., 2003: Hurricane track prediction using a statistical ensemble of numerical models. *Mon. Wea. Rev.*, 131, 749–770.
- Weisstein, E. W., 1999: Mathworld. Wolfram Web Resource, <http://mathworld.wolfram.com>.
- Whitaker, J. S. and Hamill, T. M., 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, 130, 1913–1924.
- Wikipedia, 2001: On-Line Encyclopedia, <http://en.wikipedia.org>.
- Wilson, J. W., 2004: Precipitation nowcasting: Past, present and future. *Sixth International Symposium on Hydrological Applications of Weather Radar*, Australian Government Bureau of Meteorology, keynote address, available on-line or on conference proceedings CD. [http://www.bom.gov.au/bmrc/basic/contents/qpf/WILSON\\_KEYNOTE.pdf](http://www.bom.gov.au/bmrc/basic/contents/qpf/WILSON_KEYNOTE.pdf).
- Wunsch, C., 1996: *The Ocean Circulation Inverse Problem*. Cambridge University Press, 442pp.

- Yeung, P. K., 2002: Lagrangian investigations of turbulence. *Ann. Rev. Fluid Mech.*, 34, 115–142.
- Zhang, F., 2004: Dynamics and structure of mesoscale error covariance in the snowstorm of 24–25 January 2000 estimated through short-range ensemble forecasts. *Mon. Wea. Rev.*, submitted.
- Zhang, F., Snyder, C., and Rotunno, R., 2002: Mesoscale predictability of the "surprise" snowstorm of 24–25 January 2000. *Mon. Wea. Rev.*, 130, 1617–1632.
- Zhu, L., Haker, S., and Tannenbaum, A., 2003: Area-preserving mappings for the visualization of medical structures. R. E. Ellis and T. M. Peters, eds., *MICCAI 2003*, Springer-Verlag, 277–284.
- Zou, X. and Xiao, Q., 2000: Studies on the initialization and simulation of a mature hurricane using a variational bogus data assimilation scheme. *J. Atmos. Sci.*, 57, 836–860.
- Zwiers, F. W. and von Storch, H., 1995: Taking serial correlation into account in tests of the mean. *J. Climate*, 8, 336–351.