Abstract

FOLEY, KRISTEN MADSEN. Multivariate Spatial Temporal Statistical Models for Applications in Coastal Ocean Prediction. (Under the direction of Montserrat Fuentes and Lian Xie.)

Estimating the spatial and temporal variation of surface wind fields plays an important role in modeling atmospheric and oceanic processes. This is particularly true for hurricane forecasting, where numerical ocean models are used to predict the height of the storm surge and the degree of coastal flooding. We use multivariate spatial-temporal statistical methods to improve coastal storm surge prediction using disparate sources of observation data. An Ensemble Kalman Filter is used to assimilate water elevation into a three dimension primitive equations ocean model. We find that data assimilation is able to improve the estimates for water elevation for a case study of Hurricane Charley of 2004. In addition we investigate the impact of inaccuracies in the wind field inputs which are the main forcing of the numerical model in storm surge applications. A new multivariate spatial statistical framework is developed to improve the estimation of these wind inputs. A spatial linear model of coregionalization (LMC) is used to account for the cross-dependency between the two orthogonal wind components. A Bayesian approach is used for estimation of the parameters of the multivariate spatial model and a physically based wind model while accounting for potential additive and multiplicative bias in the observed wind data. This spatial model consistently improves parameter estimation and prediction for surface wind data for the Hurricane Charley case study when compared to the original physical wind model. These methods are also shown to improve storm surge estimates when used as the forcing fields for the coastal ocean model. Finally we describe a new framework for estimating multivariate nonstationary spatial-temporal processes based on an extension of the LMC model. We compare this approach to other multivariate spatial models and describe an application to surface wind fields from Hurricane Floyd of 1999.

Multivariate Spatial Temporal Statistical Models for Applications in Coastal Ocean Prediction

by

Kristen M. Foley

a dissertation submitted to the graduate faculty of
north carolina state university
in partial fulfillment of the
requirements for the degree of
doctor of philosophy

statistics

raleigh
2006

approved by:

Dr. Montserrat Fuentes (Co-Chair)

Dr. Lian Xie (Co-Chair)

Dr. Jerry Davis

Dr. David Dickey

Dr. Sujit Ghosh

Dedication

To my husband, Derrick Foley.

Biography

Kristen Foley was born on April 10, 1979, in Charlotte, North Carolina. In 1997 she graduated from Northwest Cabarrus High School in Concord, N. C. She attended N. C. State University to study engineering but was thankfully recruited to the statistics program in her freshman year. She decided to stay in the program for her master's degree and eventually her PhD, which would prove to be an excellent decision. She began working with Dr. Montserrat Fuentes during her first year on problems in environmental statistics. She collaborated on projects with the U.S. Environmental Protection Agency to improve air quality maps and with the NCSU Coastal Fluid Dynamics Laboratory to improve storm surge prediction. During her time in the graduate program she worked with other students to organize a day–long Career Fair and a new Teaching Working Group and a Research Working Group. She was also involved in several social activities including organizing a Departmental Feud Game Show, a canoe trip and multiple appearances in the annual International Dinner talent show. She begins work at the EPA in September.

Acknowledgements

This study is supported by the Carolina Coastal Ocean Observation and Prediction (Caro–COOPS) program funded by the National Oceanic and Atmospheric Administration (NOAA) through the NOAA Coastal Services Center in Charleston, South Carolina. Additional funding was provided by the NCSU Department of Statistics through a grant from the National Science Foundation. I greatly appreciate the help and advice on the use of the ocean model application provided by members of the Coastal Fluid Dynamics Laboratory headed by Dr. Lian Xie at North Carolina State University. The feedback and advice provided by Dr. Lian Xie, Dr. Sujit Ghosh, Dr. Jerry Davis and Dr. Dave Dickey were invaluable in completing this dissertation. Dr. Montserrat Fuentes provided support and guidance throughout my entire five years in the Statistics graduate program. I am tremendously grateful for her mentoring and instruction throughout that time.

My appreciation also goes to Dr. William Swallow who first recruited me to the NC State Statistics program and could always be trusted to have a word of encouragement and some sound advice whenever I wandered into his office. I was blessed to have a wonderful group of friends in the department that quickly become my extended family as we struggled through homework sets and review sessions together. I would never have enjoyed the graduate program so much if it were not for Lovely, Alvin, Marti, Lavanya, Karen, Joe and many others.

Above all I would like to acknowledge the profound influence of my family. My sister, Kimberly, and my husband, Derrick, have had to see me on a daily basis as I stressed over finishing this dissertation. I am grateful for their patience and their love. The support of my parents, Ron and Marty, is what gets me through all of the challenges in my life, great and small. They are the solid ground beneath my feet that allows me to push myself and dream big dreams.

Table of Contents

List of Tables

List of Figures

Chapter 1

Data Assimilation for Coastal Ocean Prediction

1.1 Introduction

Multivariate spatial statistical problems are prevalent in the environmental sciences, particularly in atmospheric and oceanic data applications. In these cases the state space is typically of very high dimension and the processes of interest are inherently nonlinear and dynamic. Different sources of information for these systems include observational data as well as numerical models which are based on the fundamental physical principles driving atmospheric and oceanic processes. Over the past decade there has been an increase in the amount of available real time observed oceanic data as well as advances in the sophistication and resolution of deterministic ocean models. Whether the state variables of interest are directly observed (e.g. buoys and ship data) or indirectly observed (e.g. satellite data), these observations may be at very different scales in space and time. Our goal is to obtain more accurate prediction of wind fields to improve the quality of storm surge forecasts for the Eastern US coastal region. These forecasts can be used for assessments of warnings and evacuation notices and can provide valuable information for recovery operations and response planning.

In order to use all available coastal observations we use data assimilation (DA) which is the process of combining noisy observations with numerical ocean model output. Multivariate spatial–temporal modeling techniques are applied to better model hurricane wind fields that combine wind information from buoys, ships, satellites, and physical models. These wind fields are used as the primary forcing for numerical forecasts of the coastal ocean response to hurricane force winds and pressure.

In the following sections of this chapter we introduce the concept of data assimilation and describe a case study for Hurricane Charley of August 2004. In Chapter 2 we propose a Bayesian framework to capture different sources of uncertainty and bias in available observed wind data and physics–based wind models. A spatial linear model of coregionalization (LMC) is used to explain variability in the horizontal and vertical wind components as well as the cross–dependency between these two components. These methods are then compared with the Ensemble Kalman Filtering methods to explore the impact on storm surge forecasts for the same case study of Hurricane Charley. In Chapter 3 we review multivariate spatial models and propose a new flexible multivariate nonstationary spatial–temporal statistical model based on an extension of the LMC approach. We conclude with an application to wind fields from Hurricane Floyd of 1999.

The following chapter is organized as follows. Section 1.2 describes the storm surge application, the available observed data and the numerical ocean model. Section 1.3 reviews the sequential Kalman Filtering methods including the Ensemble Kalman filtering technique proposed by Geir Eversen (Evensen, 1994) that is used here. Section 1.4 describes the results of a case study of Hurricane Charley of August 2004. In a series of simulations we use model generated "pseudo" observation to test the effectiveness of the EnsKF approach under two different sources of modeling error: errors in initial conditions and errors in the surface forcing fields. We also test the benefit of additional water gauge sites along the coast of Georgia and the Carolinas by using a set of existing and hypothetical networks. A concluding discussion is given in section 1.5.

1.2 Data and Scientific Problem

1.2.1 Data Assimilation for Storm Surge Forecasting

Storm surge is the abnormal rise of water in the coastal ocean, estuaries and lakes caused by the high winds and to a lesser extent the low pressure associated with a hurricane or extratropical cyclone. Storm surge can compound the effects of inland flooding caused by rainfall, leading to loss of property and loss of life for residents of coastal areas. Historically storm surge has been responsible for the largest loss of life and property damage since the effects of the flooding can last long after the hurricane has passed (Powell and Houston, 1996). In order to improve storm surge forecasts data assimilation methods are used to combine output from a numerical ocean model with coastal observations of the storm surge height zero to six hours before a hurricane makes landfall.

Clearly, numerical ocean models are essential forecasting tools for coastal areas most likely to be impacted by storm surge. These models are also used for creating "nowcast" estimates which are forecasts for short time scales typically less than one hour. However model errors are inevitable due to uncertainty in initial and boundary conditions as well as simplified or neglected physical processes or mathematical approximations used in the system equations. On the other hand, unlike numerical weather prediction applications, available coastal ocean observations are relatively sparse and may not always capture complex coastal dynamics (Echevin et al., 2000). Measurement error is also a source of variability that is often not well understood and so hard to characterize.

Data assimilation (DA) is the process of combining noisy observations with output from physics–based numerical models. The problem of data assimilation for the oceanic and atmospheric sciences fits naturally into a Bayesian framework. That is, we observe data based on the underlying (unknown) state of the atmosphere or ocean and we estimate this current state based on what has been observed. State of the art data assimilation techniques have

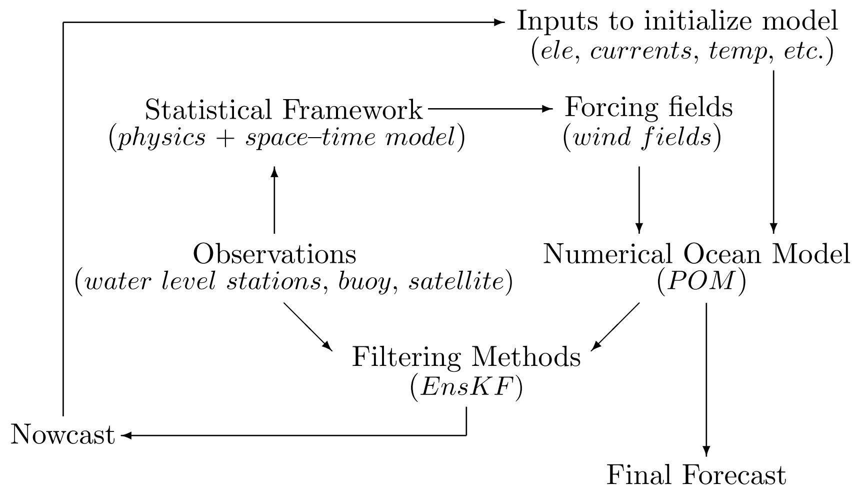

Figure 1.1: Scientific framework for coastal ocean prediction.

mainly been developed for applications in numerical weather prediction. Only in the past decade has interest begun to grow in applying these data assimilation methods to oceanic modeling. The main impetus for this new area of research is the substantial increase in the type and number of observations available for the coastal and open ocean including the expansion of permanent monitoring networks, the development of lagrangian ocean data and satellite observations. This observed data can also be used in separate stages of data assimilation as shown in the flow chart in Figure 1.1. Observations can be used to estimate input fields or optimize parameter values that are used to initialize the numerical model. This step is the topic of Chapter 2. The focus of this chapter is on the filtering methods that combine model output with observed data.

In the 1980s and early 1990s several studies were done to test different oceanic DA schemes in relatively simple models with limited physics. More recently work has been done using primitive equation models with realistic topography. In a series of papers, Mellor and Ezer described the assimilation of satellite–derived sea surface data into the Princeton Ocean Model (POM), a three dimensional ocean model, using an optimal interpolation (OI) scheme (Mellor and Ezer (1991), Ezer and Mellor (1994), Ezer and Mellor (1997)). The OI method uses a predetermined and fixed (in time) error covariance. Implications of using a fixed error covariance will be discussed in section 1.3.

Over the past decade more advanced DA methods have been tested for variety of applications. Evensen and van Leeuwen (1996) used the Ensemble Kalman filter (EnsKF) to assimilate satellite data into a two layer quasi–geostrophic model for a region off the coast of South Africa. Verlaan and Heemink (1997) proposed a Reduced Rank Square Root Kalman Filter (RRSQRT KF) to do tidal flow forecasting using a two dimensional shallow water equation model. Madsen and Canizares (1999) compared results from the RRSQRT KF and EnsKF using simulated surface elevation observations (i.e. model generated data plus noise). They found the RRSQRT KF was more efficient in correcting model errors except in the case of strongly non–linear dynamics. Echevin et al. (2000) used POM under an idealized setup with a baroclinic current flowing onto a coastal shelf. They compared two OI methods with an Ensemble Kalman Fitler (EnsKF) and showed how the impact of assimilated measurements on the analysis is dependent on the location of the observations relative to the coastline and topography. Luong et al. (1998) used a 4D variational approach to assimilate satellite derived sea–surface heights into a quasi–geostrophic ocean circulation model.

In terms of data assimilation for storm surge forecasts, Heemink et al. (1997) looked at forecasts using the KF and the RRSQRT KF. Canizares et al. (1996) and Canizares et al. (1998) used a 2D shallow water model to apply the RSQRT KF with correlated system noise. They compared the DA results for a 1993 storm using observations from 9 stations often impacted by storm surge along the coast of the North Sea.

In this chapter, a sequential data assimilation approach is utilized to improve the quality of storm surge nowcasts and forecasts for the Eastern US coastal region. These storm surge estimates can be used for assessments of warnings and evacuation notices and can provide valuable information for recovery operations and response planning. This is one of many applications in the physical sciences where reliable estimation of the state variables of interest relies on properly identifying and quantifying disparate sources of uncertainty in the data analysis. For example, it is often unclear whether the main source of prediction errors is simplified or neglected physical processes used in the numerical model, inaccuracies in the initial conditions or errors in the forcing fields. Prediction errors can certainly arise from a combination of these factor but the challenge is to decide how to allocate time and computing resources.

Our ultimate goal is to develop methods that would be practical (in terms of computing time) for real–time forecasting. The sophistication of any proposed methodology must be constrained by the computing demands of the method. We explore what will have the largest impact on real–time estimation of coastal water levels under hurricane forcings.

1.2.2 Princeton Ocean Model

The Princeton Ocean Model (POM) is used for coastal ocean modeling, which is adopted by the National Oceanic and Atmospheric Administration (NOAA) as the operational coastal ocean forecasting system for the U. S. East Coast. POM is a fully three–dimensional numerical model and has the capability of modeling the evolution of ocean processes at many layers. In addition, POM uses a bottom following sigma coordinate system to account for differences in water depth. As a result, the POM model is able to predict storm surge with greater accuracy than traditional depth–averaged storm surge models (Peng et al. (2004); Pietrafesa et al. (1997)).

Specifying the bathymetry of the coastal region being studied plays an important part in storm surge forecasting. According to the Federal Emergency Management Agency: “The level of surge in a particular area is also determined by the slope of the continental shelf. A shallow slope off the coast will allow a greater surge to inundate coastal communities. Communities with a steeper continental shelf will not see as much surge inundation, although large breaking waves can still present major problems.” Thus applying the ocean model at a higher resolution to capture such smaller scale phenomena is essential for improving the accuracy of coastal forecasts. POM has been used in various applications to model estuaries and bays, semi–enclosed seas and coastal regions around the world and has been applied extensively for studies in the Gulf Stream region (see Ezer and Mellor (1997) and references therein). It has been designed to represent ocean phenomena of 1–100km length and tidal–monthly time scales. For more details on the governing equations and parameterizations used in the Princeton Ocean Model see Blumberg and Mellor (1987).

For the case study presented here POM is used to simulate the coastal ocean response to strong hurricane forcing conditions along the coast of the Carolinas and Georgia. The ocean model is run at a resolution of 1 minute longitude by 1 minute latitude ($\approx$ 1.5 km by 1.9 km) using an inundation model developed by Xie et al. (2004) and applied to the Pamlico Sound system by Peng et al. (2004), that allows for flooding and drying in shallow coastal areas. Bathymetry for this gridded domain is obtained from the National Geophysical Data Center's (NGDC) Coastal Relief Model as shown in Figure 1.2.

1.2.3 Real Time Water Level Observations

Observed water levels were used from a series of coastal water level stations maintained by NOAA's National Ocean Service (NOS) Center for Operational Oceanographic Products and Services (CO–OPS) as well as those maintained by the Carolina Coastal Observation and Prediction Program (Caro–COOPS). The state variable of interest is the *change* in surface water elevation from an initial value of zero meters which is used for storm surge forecasting. Hourly data is obtained from 7 coastal sites for Hurricane Charley based on the reported meters above mean lower low water (MLLW). The MLLW is the average of the lower low water height of each tidal day observed over a period of 19 years and is used to adjust for differences in tidal gauge locations.

Since these observed values will be influenced by the tidal signal as well as the hurri-

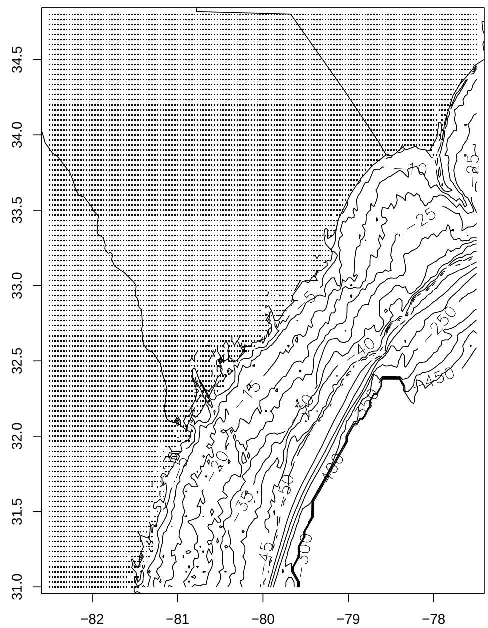

Figure 1.2: Bathymetry (in meters) for model domain.

cane wind forcing, tidal predictions are also pulled from CO–OPS website. The assimilated observations are thus the adjusted water level: meters above MLLW minus the predicted hourly tides. These adjusted observations represent the change in water elevation that can be attributed to storm surge and in this way can be compared to the state variable in the numerical model. An alternative approach would be to use a coupled tide–surge model that includes the tidal signal in the integration of the numerical model (e.g. Jones and Davies (1998)).

1.3 Coastal Ocean Data Assimilitaion

1.3.1 Kalman Filter

Data assimilation is a type of state space estimation problem where we estimate the state of a system that changes over time using noisy observations made on the system. Following the notation of Ide et al. (1997), $Y^o(t_i) = Y_i^o$ is a $(p \times 1)$ vector of observations at irregularly spaced points at time t_i and X_i^f is a $(n \times 1)$ vector consisting of the numerical model output for k model variables at j grid points in 2 or 3 dimensions at time t_i. These variables are ordered by grid point to form a single vector such that $k \times j = n$. The true unobserved state vector X_i^t is a $(n \times 1)$ vector representing the true values of the k variables discretized at the j model grid points. X_i^a will be used to denote the estimated or "analysis" state vector based on all available observations at time t_i.

The evolution of the dynamic system is described by the state or system equation:

$$X_i^t = \mathcal{M}_{i-1,i}[X_{i-1}^t] + \eta_i \tag{1.1}$$

where $\mathcal{M}_{i-1,i}[]$ represents the forward integration of the current state by the numerical model. The observation equation relates the observations to the state vector in the presence of measurement error:

$$Y_i^o = \mathcal{H}_i[X_i^t] + \epsilon_i^o \tag{1.2}$$

$\mathcal{H}_i[]$ is the observation operator which interpolates the model variables to the observation locations and transforms the model variables to the observed variables. It is typically assumed $\epsilon_i^o \sim P^o(0, R_i)$, $\eta_i \sim P^f(0, Q_i)$ and $\text{Cov}(\epsilon_i^o, \eta_i) = 0$.

In the case that ϵ_i^o and η_i are normally distributed and $\mathcal{H}_i[]$ and $\mathcal{M}_{i-1,i}[]$ are linear operators, $H_{i(n\times p)}$, $M_{i-1,i(n\times n)}$, an analytic solution exists for the posterior distribution of the true state conditioned on available data. Prior to observing Y_i^o, the numerical model represented in the state equation (1.1) is used to evolve the state vector and the error covariance. Given the estimated state vector and error covariance matrix from the last time period, (X^a_{i-1}, P^a_{i-1}), the prior distribution is:

$$P(X^t_i|Y^o_{i-1}) \sim N(X^f_i, P^f_i) \tag{1.3}$$

where

$$\begin{aligned} X^f_i &= M_{i-1,i}X^a_{i-1} \\ P^f_i &= M_{i-1,i}P^a_{i-1}M^T_{i-1,i} + Q_i \end{aligned} \tag{1.4}$$

This forecasting step is what makes the Kalman Filter so computationally expensive. The matrix multiplication by $M_{i,i+1}$ to obtain the forecast error covariance P^f_{i+1} requires approximately as much computing effort as advancing the the deterministic forecast model $n/2$ times where n, the dimension of the state vector, is typically on the order of $10^3 - 10^5$ (Kalnay, 2003).

Note that $P(Y^o_i|X^t_i, Y^o_{i-1}) = P(e_i|X^t_i, Y^o_{i-1})$, where e_i is the difference between the observation and the first guess obtained by applying the observation operator H_i to the current forecast, $e_i = Y^o_i - H_iX^f_i$. This is known as the observation increment. Since H_i, $M_{i-1,i}$, and X^a_{i-1} are known, observing Y^o_i is equivalent to observing e_i. Hence the likelihood is

$$P(e_i|X^t_i, Y^o_{i-1}) \sim N(H_i(X^t_i - X^f_i), R_i) \tag{1.5}$$

An estimate of the current state of the atmosphere based on the observation and state equations described above can be easily obtained by a direct application of Bayes Theorem (Meinhold and Singpurwalla, 1983). Given the prior distribution $P(X^t_i|Y^o_{i-1})$ and the likelihood $L(X^t_i|Y^o_i) = P(Y^o_i|X^t_i, Y^o_{i-1})$ then the posterior conditional distribution best describing the state of the atmosphere at time t_i based on all available data is given by,

$$P(X^t_i|Y^o_i, Y^o_{i-1}) = \frac{P(X^t_i|Y^o_{i-1}) \times P(Y^o_i|X^t_i, Y^o_{i-1})}{\int P(X^t_i, Y^o_i|Y^o_{i-1})dX^t_i} \tag{1.6}$$

For Gaussian prior and likelihood distributions a standard multivariate result (see Appendix A) is used to avoid the direct calculation of (1.6) to find $P(X^t_i|Y^o_i, Y^o_{i-1})$. Applying this result the joint distribution for e_i and X_i^t is:

$$\left[\begin{pmatrix} e_i \\ X_i^t \end{pmatrix} \middle| Y_{i-1}^o \right] \sim N\left[\begin{pmatrix} 0 \\ M_i X_{i-1}^a \end{pmatrix}, \begin{pmatrix} R_i + H_i P_i^f H_i^T & H_i P_i^f \\ P_i^f H_i^T & P_i^f \end{pmatrix} \right] \tag{1.7}$$

Thus the posterior conditional normal distribution is:

$$P(X_i^t | e_i, Y_{i-1}^o) = (X_i^t | Y_i^o, Y_{i-1}^o) \sim N(M_i X_{i-1}^a + K_i(Y_i^o - H_i X_i^f), (I - K_i H_i) P_i^f) \tag{1.8}$$

where K_i is the Kalman gain matrix $K_i = P_i^f H_i^T (R_i + H_i P_i^f H_i^T)^{-1}$.

Now the updated estimate for the current state vector (the analysis vector) is simply the posterior mean:

$$X_i^a = X_i^f + K_i(Y_i^o - H_i X_i^f) \tag{1.9}$$

and the updated forecast error covariance matrix equals the posterior covariance matrix:

$$P_i^a = (I - K_i H_i) P_i^f \tag{1.10}$$

This pair, (X_i^a, P_i^a) is then used in the prior distribution (1.3) to initiate the next forecasting/updating cycle for time t_{i+1}. Notice that the Kalman gain matrix will be larger when the forecast error covariance P_i^f is large compared to the observation error covariance R_i. In turn, the larger K_i matrix in (1.9) will mean a larger correction to the "first guess" vector X_i^f.

In data assimilation applications in the geophysical sciences the state vector of interest is of very high dimension and the dynamic systems being modeled are typically nonlinear. Nonlinear processes are especially important for coastal dynamics due to the complexities of local bathymetry and coastlines and the influence of wind forcings, tides, and storm surge. As a result the numerical model $\mathcal{M}_{i-1,i}$ is typically a nonlinear function of the state variables and the normality assumption for the error terms may not be realistic. The Extended Kalman Filter (ExKF) was developed for this case and involves the linearization of the dynamical and observation operators. The ExKF has historically been found to be effective for weakly non-linear systems but the computational demands of propagating the error covariance forward in time, seen in the KF, still remain.

It is relevant to note that Optimal Interpolation, mentioned in section 1.2 uses the same updating formulas as the KF. The difference in the two methods is that the OI scheme assumes the background error covariance is fixed in time so that $P_i^a = P^a$ is estimated once prior to the DA cycle rather than evolving in time as in the Kalman filtering methods.

1.3.2 Variational Methods

Kalman filter approaches are known as a type of sequential method due to the fact that if measurement error is assumed to be uncorrelated, it is possible to assimilate observations one at a time as they become available. Briefly we mention an alternative type of approach to data assimilation, known as variational methods. Variational methods build from Bayesian optimal control theory and utilize a cost (or risk) function, J(X), to minimize the distance between the model output and a set of measurements within a given time interval. Specifically, the 3D–Var approach finds the value of the state vector X that minimizes the distance to the background state and to the observations (each weighted by their inverse error covariances) using a quadratic cost function:

$$2J(X) = (X - X^b)^T B^{-1}(X - X^b) + (Y^o - H(X))^T R^{-1}(Y^o - H(X)) \tag{1.11}$$

In fact, for Gaussian model and observation errors, the solution to this minimization problem is precisely the maximum likelihood estimator. Thus it can be shown (e.g. Kalnay (2003), pg 171) that Optimal Interpolation (a sequential method) and 3D–Var are solving the same problem.

As in optimal interpolation, 3D–Var methods use a predefined error covariance that remains fixed over time. Madsen and Canizares (1999) note that this simplifying assumption is particularly unrealistic in coastal ocean modeling where the dynamics of the system are "strongly influenced by land–sea boundaries and the flooding and drying of tidal areas". This is why more advanced methods such as the Kalman filter and it variants, are much more attractive for oceanic applications.

A more advanced variational method, 4D–Var, also accounts for the dynamic propagation of model errors. J(X) in 4D–Var is a function of the distance to the background at the start of the assimilation window, i.e. $X(t_0) - X^b(t_0)$. The minimization of the cost function requires a linear tangent model which is the linearized Jocobian ($L(t_0, t_i) = \partial\mathcal{M}_i/\partial X_0$) and its adjoint ($L^T(t_i, t_{i-1})$) which takes gradient information backwards in time. The solution to 4D–Var is a new "optimized" initial state X_0 which is integrated forward by the model to form the analysis at time t_i: $X^a(t_i) = \mathcal{M}_{0,i}[X_0]$. 4D–Var does not require the model to be a linear system since it utilizes the linear tangent and adjoint models but is does assume that the dynamical model is perfect.

4D–Var is typically much more computationally expensive than the Ensemble Kalman filter methods. It also requires the development of the linear tangent and adjoint models which will be unique for each numerical model. In the realm of numerical weather prediction it is still unclear which method performs better for forecasting. In oceanic data assimilation the comparisons between these two methods are still in their very early stages.

1.3.3 Ensemble Kalman Filter

A common adaptation of the KF is the Ensemble Kalman Filter (EnsKF) proposed by Geir Eversen in 1994 which uses a Monte Carlo approximation of the forecast distribution. The Ensemble Kalman Filter approach is used to generate samples of the state vector and carry out an ensemble of data assimilation cycles. This ensemble is then used to estimate the analysis mean and error covariance. The EnsKF uses the following algorithm where the superscript (j) is used to refer to j^{th} ensemble member (Evensen, 2003).

Forecast Step:

1. Sample $X_{i-1}^{(j)} \sim P(X_{i-1}^a, P_{i-1}^a)$, $j = 1, ..., N_s$

2. Compute $X_i^{f(j)} = \mathcal{M}_{i-1,i} X_{i-1}^{(j)}$, $j = 1, ..., N_s$

3. Calculate sample mean and covariance:

$$\begin{aligned} \hat{X}_i^f &= \frac{1}{N_s} \sum_{j=1}^{N_s} X_i^{f(j)} \\ \hat{P}_i^f &= \frac{1}{N_s - 1} \sum_{j=1}^{N_s} [X_i^{f(j)} - \hat{X}_i^f][X_i^{f(j)} - \hat{X}_i^f]^T \end{aligned} \tag{1.12}$$

Analysis Step:

1. Perturb the observations at time t_i by adding random noise (typically Gaussian): $Y_i^{*(j)} = Y_i^o + \eta_i^{(j)}$, $\qquad \eta_i^{(j)} \sim N(0, R_i)$, $j = 1, ..., N_s$

2. Update each member of the forecast sample with the KF analysis step:
   $X_i^{(j)} = X_i^{f(j)} + K_i(Y_i^{*(j)} - HX_i^{f(j)})$, $j = 1, ..., N_s$
   $K_i = \hat{P}_i^f H^T (R_i + H\hat{P}_i^f H^T)^{-1}$

3. This now serves as a sample for time t_i and the process repeats.

Typically ensembles of only 10 to 100 members are used to approximate the first two moments of the update distribution. The EnsKF requires considerably less computational cost compared to the Extended Kalman Filter and it avoids problems associated with the linearization of the forecast model. It has been found that these sample estimates tend to underestimate the forecast error covariance. It can be shown that for the case of linear dynamic and observation operators and Gaussian errors for large ensemble size, N_s, the update $X_i^{(j)}$, $j = 1, ..., N_s$, is a sample from the posterior update distribution from the standard KF given in (1.8). However since the analysis step is based on the standard Kalman filter update formulas derived from these assumptions there is no theoretical justification that the updated sample from the EnsKF will be representative of the true state for nonGaussian error statistics. On the other hand, the EnsKF clearly has an advantage over OI and 3D Var methods that use a constant error covariance since it allows for the nonlinear evolution of the error statistics.

1.4 Ocean Model Application: Hurricane Charley

1.4.1 An EnsKF for Assimilation of Water Level Observations

In this section the Ensemble Kalman Filter is applied to a case study of Hurricane Charley of 2004. This case study was chosen because the hurricane surface winds induced considerable storm surge within the CaroCOOPS observing array off the coast of Georgia and the Carolinas. Storm surge predictions based on a control run of the ocean model are compared to the analysis output from the EnsKF and to the observed data.

Hurricane Charley crossed over the central Florida peninsula and moved offshore early on the morning August 14th, 2004. Charley made landfall at Cape Romain, SC at 1400 UTC (Coordinated Universal Time) on the 14th as a weakening Category 1 storm with highest winds around 80 miles per hour (70 kt/hr). It moved off shore again and then made landfall a few hours later at North Myrtle Beach, SC. Water level gauges along the coast of Georgia and the Carolinas reported storm surge heights up to 1.5 meters above normal tidal levels. The Tropical Prediction Center of the NHC reports that Charley caused insured damages of 25 million dollars in North Carolina and 20 million dollars in South Carolina.

An axis–symmetric wind model (Holland, 1980) is used to initialize POM based on track information available from NOAA for the center location and central pressure of Hurricane Charley on August 13th – 15th. The Holland model has two forcing parameters, $Rmax$ and B. $Rmax$ is the radius of maximum sustained wind of the hurricane and B defines the shape of the pressure profile. These parameters are chosen based the size and intensity of the storm ($Rmax = 46$; $B = 1.9$) using an approach suggested by Hsu and Yan (1998). See Xie et al. (2006) and Peng et al. (2004) for further details on surface wind forcings and storm surge forecasting for coastal ocean systems. The Holland model will also be dicussed in further detail in Chapter 2.

The state variable of interest is the *change* in surface water elevation (storm surge). A statistical ensemble of initial elevation fields is generated using a Gaussian random field centered at zero meters with an exponential spatial covariance model. The ocean model is run at a resolution of 1 minute longitude by 1 minute latitude with four layers of 301×229 grid points. Due to the computational cost of using such a high resolution we use an ensemble size of twenty. The analysis/forecast cycle is initiated using a 12 hour "spin–up" from August 13, hour 900 UTC to August 14, hour 900 UTC to simulate the sea state conditions along the Eastern coast leading up to the data assimiliation window (i.e. the time period six hours prior to landfall).

From August 14th, 900 UTC to August 14th, 1800 UTC the EnsKF is used to update the state vector every hour using observed data from 7 water level stations in North Carolina, South Carolina and Georgia within the model domain. The observation error is assumed to be independent with error variance for each site based on the reported one second variability in the water level values.

Figure 1.3 shows the observed water levels at each of the 7 water gauge sites as Hurricane Charley moves along the coast. The Sunset Beach site shows the greatest storm surge of 1.67 meters. The dashed and dotted line shows the ocean model output of surface elevation based on the control run without DA. This control run tends to overestimate the increasing and decreasing elevation values and also exaggerates the effects of the hurricane on the coastal sites at Charleston, SC and Fort Pulaski, GA which exhibit little to no storm surge during this time period. The ensemble mean is used as the predicted storm surge values for these

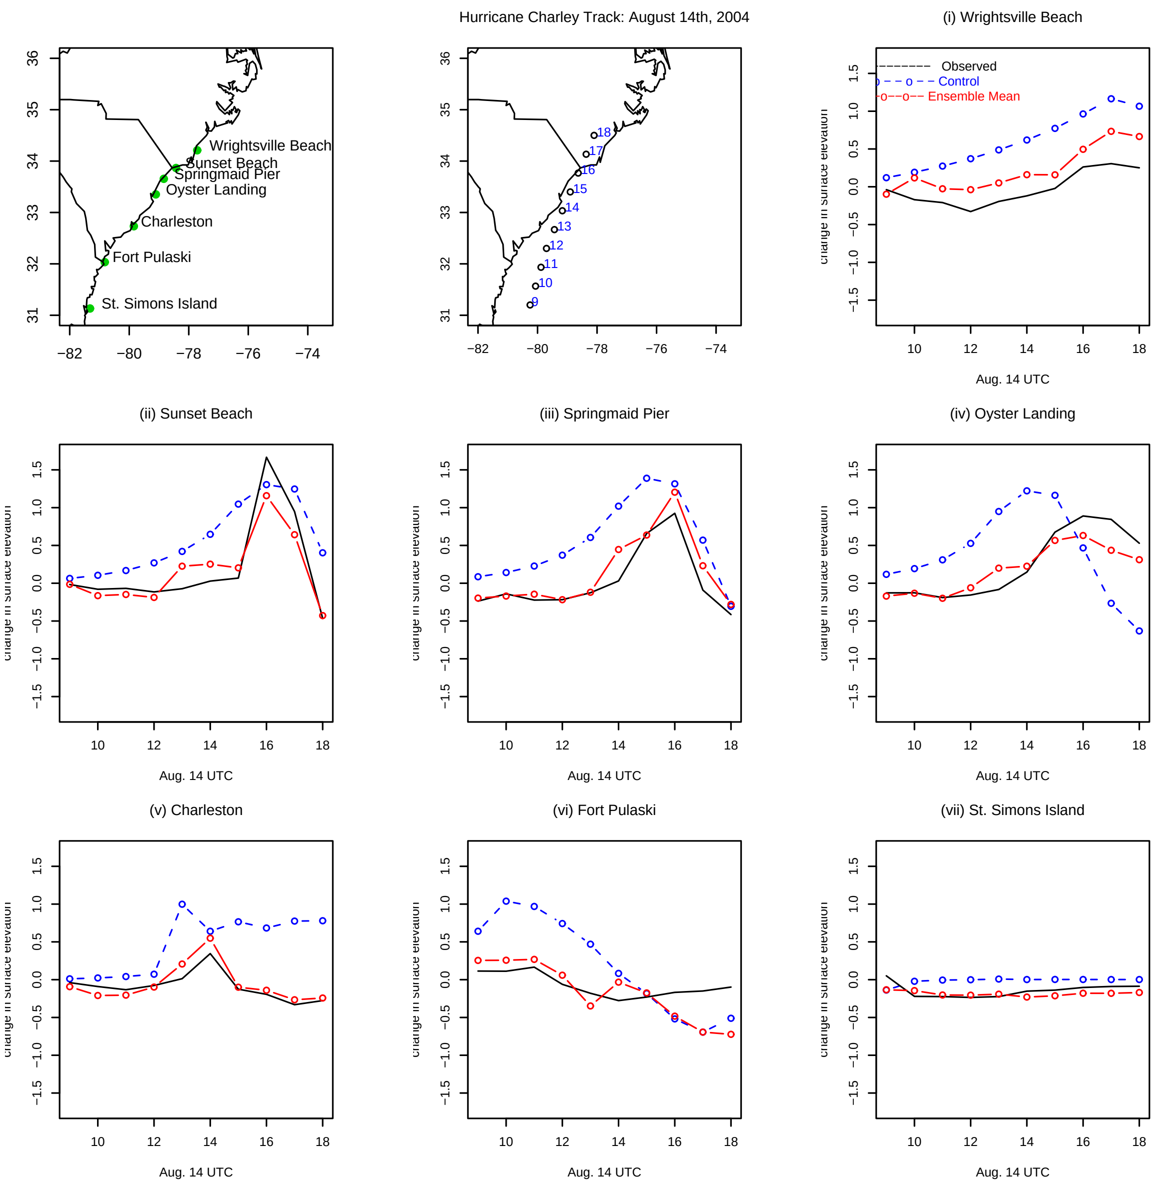

Figure 1.3: EnsKS results with the inundation model. Maps show the location of 7 coastal water elevation gauges and the track of Hurricane Charley on August 14th for hours 900 UTC to 1800 UTC. Time series plots show observed water elevation at each site as a solid black line. The dashed and dotted line shows the model output of surface elevation based on the control run without DA using the inundation ocean model. The dotted lines show the ensemble mean based on data assimilation using the EnsKF.

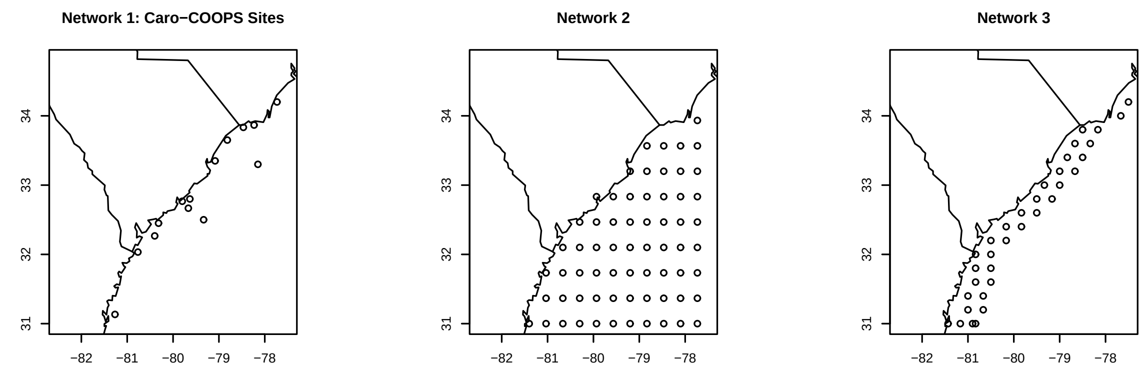

Figure 1.4: Hypothetical networks 1, 2 and 3 in ocean model domain. Network 1 includes existing locations of Caro–COOPS and CO–OPS moorings and water level sites.

times and is plotted as a dotted line. The EnsKF adjusts the model output and reduces prediction error when compared to the control run. The ensemble mean better captures the magnitude and the timing of the peak surge at these sites. This improvement is most notable at the Oyster Landing and Charleston sites. However both control run and EnsKF analysis output underestimate the largest surge value at the Sunset Beach location.

1.4.2 EnsKF Twin Experiments

We now investigate whether errors in initial conditions or in wind forcing parameters play a larger role in hurricane storm surge forecasting. The Ensemble Kalman Filter adjusts the estimate of the state vector X_i^t which is then used as the new initial conditions for the subsequent forecast. If model errors are more influenced by errors in the forcing fields we may wish to concentrate computing time on improving estimates of the surface wind fields rather than correcting errors in the initial conditions. We conduct a series of identical twin experiments meaning the model is assumed to be perfect. “Psuedo” observations are used for the data assimilation step, i.e. model generated data plus simulated measurement error. Then the analysis fields are compared to the “true” values from the original model output.

In the first set of experiments the true state is defined as model generated output for hours 900 UTC to 1700 UTC, August 14th, based on a model run with a 12 hour spin–up from August 13 hour 900 UTC to August 14th hour 9 UTC. The control run is then model output for this 9 hour period based on a model run with no spin–up. In other words, the "true" initial elevation fields for hour 900 UTC are simulated by POM based on 12 hours of hurricane force winds, but our control run assumes that the initial elevation fields are uniformly zero at this hour. Under this set–up, the impact of adding additional water level gauges within the model domain of interest is also tested. Figure 1.4 depicts three monitoring networks along the coast of the Carolinas and Georgia. Network 1 includes existing CO–OPS locations as well as moorings and water level sites from the Carolinas Coastal Ocean Observing and Prediction System (Caro–COOPS). The remaining two hypothetical networks are used to study the impact of additional sites by comparing the difference in adding these sites evenly throughout the domain (Network 2) versus concentrating the new sites along the coast (Network 3). It is not expected that such a large number of extra sites would ever realistically be used but this exaggerated design is used to help determine the benefit of additional water gauge sites and to evaluate the effectiveness of the current network.

Based on pseudo observations from the true model run at each observation network the Ensemble Kalman filter is used to assimilate the noisy (i.e. including measurement error) hourly water elevation values at hours 900 UTC to 1400 UTC. The model is then integrated forward without the benefit of DA for a three hour forecast from 1500 UTC to 1700 UTC. Since we are typically most interested in the areas along the coast with the greatest amount of storm surge we compare the prediction errors (observed water level – predicted water level) at only 10% of the surface grid points where the largest increase in the true elevation values is observed (i.e. there are a total of 17365 surface layer grid points and we consider the errors at 1736 locations near the coast).

Figure 1.5(a) shows boxplots of these prediction errors for the control run and each of

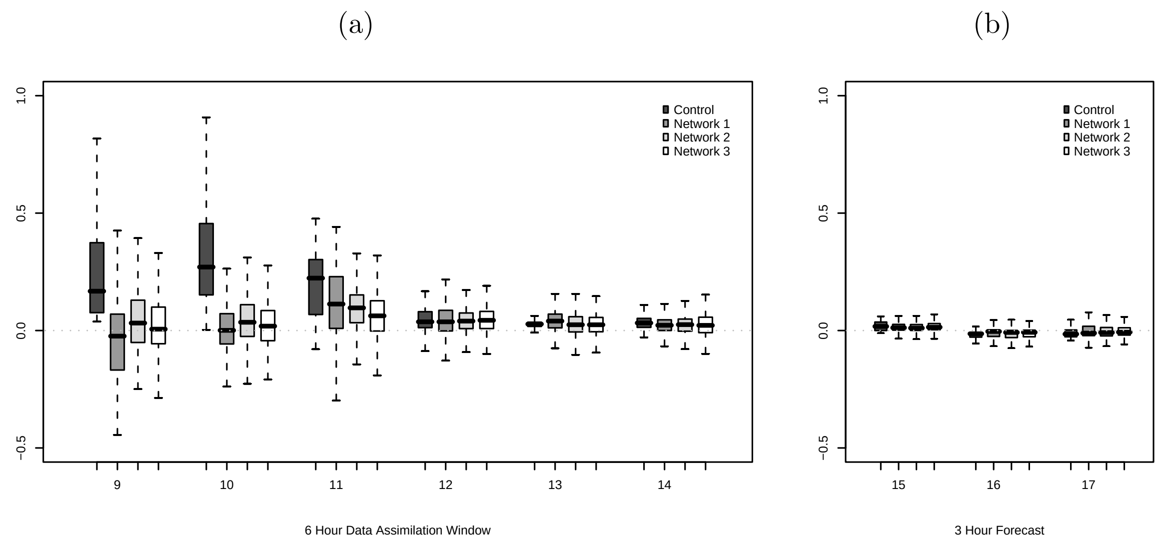

Figure 1.5: Boxplots for prediction errors (in meters) under Experiment 1 (a) for DA cycle August 14 900 UTC to 1400 UTC, (b) for 3 hour forecast from 1500 UTC to 1700 UTC.

the three network designs for the six hour DA cycle. At the beginning of the DA cycle there are prediction errors up to 1 meter in the control run. During this time the control run exhibits large bias (the median bias is depicted by the bar in the center of each boxplot). The DA estimates are less biased and show less error variance as well. Although there is some reduction in the model errors based on the DA using the different network designs, this improvement is only evident at the beginning of the DA window. Even Network 2 which has 65 locations does not greatly reduce prediction errors after the third hour when compared to the Network 1 with only 14 sites. The errors in the control and the DA output decrease over time. Both the true model run and the control run used the same wind forcing fields. Errors in the initial elevation does not appear to have a lasting impact on the prediction of surface elevation in this scenario. For the three hour forecast shown in Figure 1.5(b) there is little to no remaining effect of using the incorrect initial conditions at hour 900 UTC and all of the forecast values are very close to the true elevation at these times.

For the second set of experiments the true state for hours 900 UTC to 1700 UTC, August

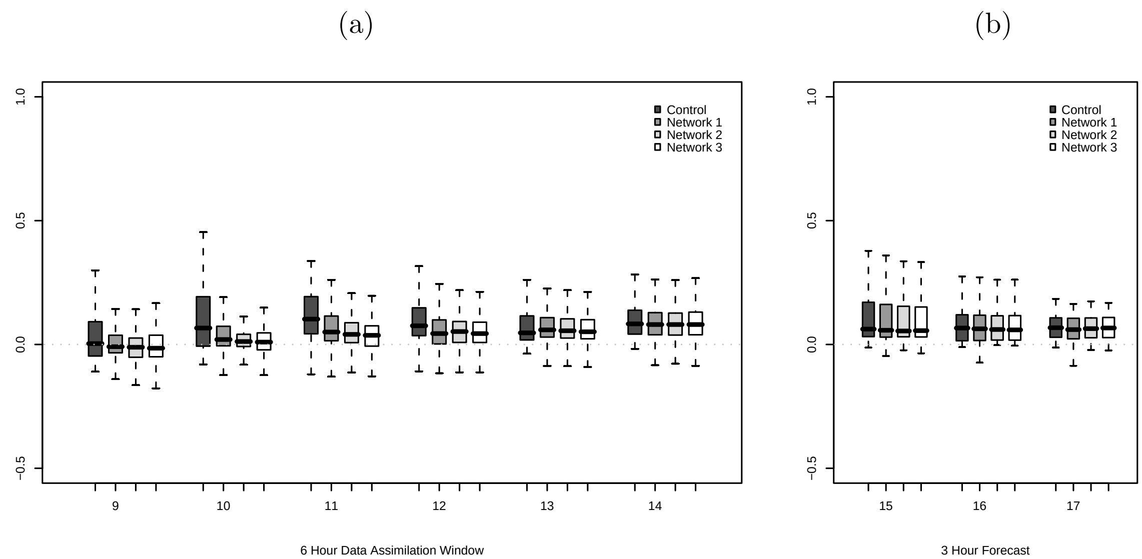

Figure 1.6: Boxplots for prediction errors (in meters) under Experiment 2 (a) for DA cycle August 14 900 UTC to 1400 UTC, (b) for 3 hour forecast from 1500 UTC to 1700 UTC.

14th, is defined as model output based on a model run using a 12 hour spin–up under forcing fields with $Rmax = 46$ and $B = 1.9$. For the control run we treat these "true" parameters as unknown and fix $Rmax = 35$ and $B = 1.0$ indicating a much weaker storm. We repeat the DA/Forecast cycles as before for the three observation networks. The boxplots of the prediction errors are given in Figures 1.6(a) and (b). In this experiment we see the model error increases over time as the hurricane eye moves into the model domain and then decreases at hour 17 UTC as Charley moves inland and the surface wind induced storm surge is greatly reduced.

Even though the EnsKF is correcting initial conditions not errors in the forcing fields, these plots show some improvement in prediction during the six hour DA window. The improvement is much less during the three hour forecast cycle. Again there is not a significant difference between the results based on the three different observation networks. In general the additional observation locations do not have much impact on reducing model error when compared to the network of current locations (Network 1). There is a consistent bias in the estimates in the control run. As expected the predicted values tend to be smaller than the true values due to the incorrect specification of the wind forcing fields. As this error in the forcing fields persists over time the data assimilation results do not fully correct this bias. This is especially true for the forecasting errors which are almost all positive for the control run as well as the DA estimates. This suggests that improvement in the wind forcing fields are needed to achieve improved forecasts. As shown in Figure 1, wind observations can also be used in the DA cycle through inverse modeling techniques or statistical modeling to improve the method of calculating these input fields. This issue is further explored in Chapter 2.

1.5 Discussion

In this chapter we considered data assimilation results for real–time estimates or nowcasts of hurricane induced storm surge. We found we were able to improve the output from the numerical ocean model at times when observational data is available. A separate but very important issue is the impact of data assimilation on forecasts. In this case the data assimilation results are used as improved initial conditions when the ocean model is integrated forward. Our analysis of these forecasting results through the identical twin experiments show that the forecast ensemble mean follows the control run after only a few hours. This suggests information from the observed water level is propagated to the state variables of the numerical model only during the assimiliation window. As a result the forecast storm surge predictions do not show an improvement over the control run except for short term nowcasts.

We test the effect of the choice of the ensemble size and the values of the error covariance used to create the initial ensemble of perturbation fields. Using a lower resolution ocean model we run the data assimiliation analysis with 80 ensembles as well as 20 and find very little difference in the ensemble mean values used for prediction. We also vary the spatial correlation structure of the error covariance defined by the Gaussian random fields by using different values for the parameters of the exponential covariance formula. We find these changes also have very little impact on the final analysis fields produced by the Ensemble Kalman filter. Consistently in each data assimilation run the variability or “spread” of the ensembles quickly decreases during the several hours of forecasting/updating steps. This is due in large part to the influence of the surface wind forcings driving the ocean model which is the topic of Chapter 2. In the following chapter we describe a methodology for including disparate sources of surface wind field data into the data assimilation scheme for coastal ocean prediction. We shall see that we can also obtain improved storm surge estimates by improving the wind field inputs.

Chapter 2

A Bayesian Framework for Multivariate Spatial Wind Field Prediction

2.1 Introduction

Estimating the spatial and temporal variation of surface wind stress fields plays an important role in modeling atmospheric and oceanic processes. This is particularly true for hurricane forecasting where numerical ocean models are used to simulate the coastal ocean response to the high winds and low pressure associated with hurricanes, such as the height of the storm surge and the degree of coastal flooding. According to the National Oceanic and Atmospheric Administration's (NOAA) Hurricane Research Division more than eighty–five percent of storm surge is caused by winds pushing the ocean surface ahead of the storm. The numerical ocean models used for storm surge forecasting are also driven primarily by the surface wind forcings. Houston and Powell (1994) did an analysis of the impact of wind field forcings on the storm surge forecasts used by the National Hurricane Center (NHC). They concluded that up to 6 hours before landfall, real–time model runs can be used to evaluate warnings and assess the extent of storm surge inundation. Houston et al. (1999) state that "an accurate diagnosis of storm surge flooding, based on the actual track and wind fields could be supplied to emergency management agencies, government officials, and utilities to help with damage assessment and recovery efforts." Our objective is to use all available sources of wind data in order to improve the estimation of the wind field inputs. It is expected that improving these inputs will have a great impact on real–time storm surge estimates, allowing recovery efforts to be organized immediately after a hurricane has made landfall according to the areas most impacted by the storm.

Several studies have investigated the potential of using real–time observation–based winds to improve storm surge predictions as new wind observations have become available from the Hurricane Research Division (HRD). Since 1993 the HRD has reported real–time analyses of tropical cyclone surface wind observations for evaluation and eventually to be utilized by the NHC (Houston and Powell, 1994). Powell and Houston (1996) did a study of Hurricane Andrew (1996), and concluded that real–time wind fields would improve storm surge forecasts, particularly for storms with asymmetric (i.e. non axis–symmetric) wind fields. Houston et al. (1999) used the HRD winds as input into a two dimensional storm surge model. In their analysis they did not allow the wind and surface pressure field to change over time, but rather ran the ocean model assuming the wind and surface pressure fields were steady state for over 24 hours.

We propose a new method to use the HRD data to fit a statistical model for the winds that can be used as input into a numerical ocean model. The statistically–based winds are a function of the hurricane's location and can be updated at intervals as small as 10 minutes as more observations become available. A statistical framework is used to account for systematic and random errors in the observations while using these observations to model the surface winds as a multivariate spatial process. This is the first time that these methods have been applied to an ocean model application to be used for improved storm surge forecasting.

The statistical modeling of spatial temporal processes in the atmospheric sciences has been an active area of research over the past decade. Hierarchical Bayesian methods have been found to be particularly useful for such applications, e.g. see Wikle et al. (1998), Berliner et al. (1999) and Royle et al. (1999). Such methods have been applied to modeling wind fields. Wikle et al. (2001) present a hierarchical spatial–temporal model for surface wind fields over the tropical oceans. Fuentes et al. (2005) propose a method for predicting coastal wind fields based on a nonstationary and nonseparable covariance model that accounts for bias in available observations. In both of these studies the orthogonal components of the vector wind fields are treated as independent. Cripps et al. (2005) apply generalized autoregressive conditional heteroscedastic (GARCH) models to the bivariate wind vector at a single spatial location. In our statistical framework we combine all sources of available data including a physics–based wind model and the observational data currently available to NHC for coastal ocean prediction along the Eastern coastline. We investigate the importance of accounting for the cross–covariance of hurricane surface wind fields. The ultimate goal in the application presented here is to develop methods that would be practical (in terms of computing time) for a real–time forecast or nowcast scenario. Hence the sophistication of any proposed methodology must be constrained by the computing demands of the method. The scientific question of interest is how to balance the sophistication of the wind field model with computational cost.

Section 2.2 describes the deterministic physical model for surface wind fields, the available observed data and the proposed statistical model. Section 2.3 reviews statistical methodology for multivariate data following the linear model of coregionalization used in this paper. We outline a hierarchical Bayesian approach for estimation of parameters in a deterministic wind model and prediction of gridded surface wind fields. Section 2.4 describes the results of a case study for wind field data from Hurricane Charley in August of 2004. The estimated wind fields are used to initialize the numerical ocean model and are shown to improve storm surge prediction. A concluding discussion is given in section 2.5.

2.2 Data and Scientific Problem

2.2.1 Modeling Surface Wind Fields

For coastal ocean modeling we use the Princeton Ocean Model (POM), which is a three–dimensional primitive equations ocean model (Blumberg and Mellor, 1987). Surface wind fields are the most important input variable used to initialize the numerical ocean model to simulate the coastal ocean response to hurricanes. A wind field is made up of wind vectors defined at a finite set of grid points such that each vector has a magnitude or wind speed value (typically measured in meters per second or knots) and a direction. These vectors are decomposed into orthogonal u winds, (East–West) and, v winds, (North–South). A Bayesian hierarchical framework is used to combine different sources of wind data at different locations in order to predict the wind fields at the grid locations used by the ocean model. Let $\mathbf{V}(\mathbf{s},t) = (u(\mathbf{s},t), v(\mathbf{s},t))^T$ be the underlying unobserved spatial process for the wind components at location $\mathbf{s}$ and time t. The standard Cartesian decomposition is used to define the wind vectors since the observed and predicted winds provided by NOAA are based on this decomposition. However we note that other decompositions are possible and in particular the u and v components could also be defined according to the direction of the motion of the storm at a given time.

Currently for storm surge forecasting a deterministic wind model is used, referred to as the Holland wind model (Holland, 1980) which determines the wind velocity induced by the hurricane pressure gradient under cyclostrophic balance. The Holland wind model for surface wind speed at a given location $\mathbf{s}$ and time t is:

$$W(\mathbf{s},t) = \left[\frac{B}{\rho} \left(\frac{Rmax}{r} \right)^B (P_n - P_c) \exp^{-(Rmax/r)^B} \right]^{1/2} \tag{2.1}$$

P_n is the ambient pressure, P_c is the hurricane central pressure at time t, ρ is the air density (fixed at 1.2 kg m^{-3}), and r is the distance from the hurricane center to location $\mathbf{s}$ at time t.

Under the assumption that the hurricane wind field is axis–symmetric the u and v components can be determined from the wind speed:

$$u^H(\mathbf{s},t) = W(\mathbf{s},t) sin\phi \tag{2.2}$$

$$v^H(\mathbf{s},t) = W(\mathbf{s},t) cos\phi \tag{2.3}$$

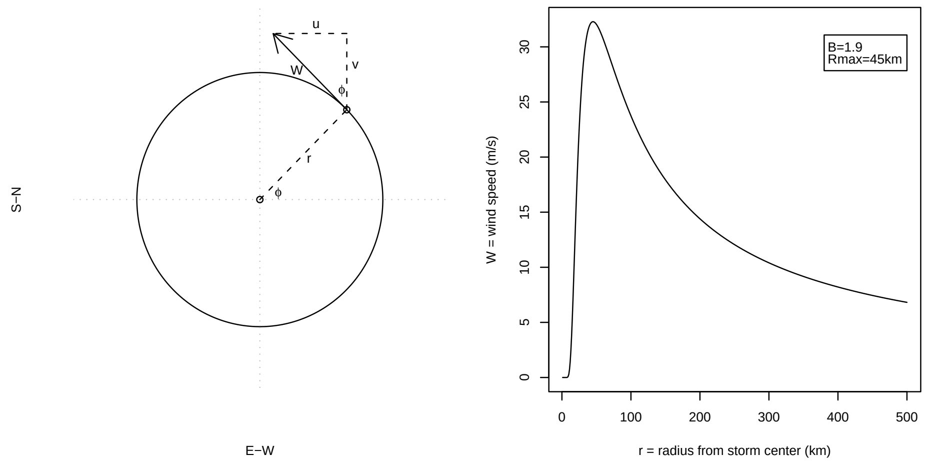

Figure 2.1: U–V wind decomposition for symmetric Holland wind model.

where ϕ is the inflow angle at site $\mathbf{s}$ across circular isobars toward the storm center at time t. That is, for every location the covariates for u and v are the radius (r), and angle (ϕ). The wind components are a nonlinear function of these covariates with parameters $\theta^H = (B, Rmax)$. The forcing parameter $Rmax$ is the radius of maximum sustained wind of the hurricane. In the Holland model the maximum possible wind speed in the storm is a function of the central pressure of the storm and the value B which defines the shape of the pressure profile. Large values of B correspond to a wind field with very high wind speeds in a small region near the center of the storm. A smaller B value will mean high wind speeds across a greater radius of the storm. These parameters are typically based on results from previous studies and are held constant for all forecasting time periods.

Although this physical model incorporates important information provided by the observed central pressure and the location of the eye of the storm there are known deficiencies in this formulation. For example, the Holland winds are symmetric around the storm center, however it is known that wind speeds are typically higher on the right hand side of the hurricane (with respect to the storm movement). Various adaptations of this physical model have been proposed. Xie et al. (2006) create an asymmetric model by incorporating data provided by the National Hurricane Center (NHC) guidance on the maximum radial extent of winds of a given threshold in the four quadrants of the storm. The NHC uses a parametric wind formula similar to the Holland model to force the ocean model. Another approach would be to use a coupled atmospheric–oceanic numerical model to simulate the surface winds at the boundary layer of the ocean model. Global mesoscale numerical weather forecasting models such as the Mesoscale Model (MM5) and the Weather Research and Forecasting Model (WRF) are capable of making forecasts for surface level wind fields. However these forecasts are not suitable for real–time forecasts of hurricane winds. Historically these models have been unable to accurately reproduce the intensity of hurricane force winds. More recent reports show that MM5 and WRF are able to obtain more realistic wind speeds when run at high resolution (on the order of 1 km grid size). However the CPU time required to produce these modeled winds prevents such model runs from being used in real–time applications. The best approach for regional storm surge forecasting is still using parametric hurricane models, as discussed in Xie et al. (2006). For further details on storm surge forecasting for the coastal ocean and estuarine systems see Xie et al. (2004), Peng et al. (2004), and references therein.

Here we focus on the operational methods and observational data currently available to the NHC for coastal ocean prediction along the Eastern coastline and incorporate this information into a statistical framework. In our analysis we provide a new tool to estimate the parameters of the Holland wind function. We include the Holland winds as the mean function of a spatial statistical model and incorporate a spatial covariance to account for any structures in observed hurricane surface winds not captured by the parametric wind fields. This approach to wind field modeling provides parameter estimates and standard errors as well as estimated wind fields through spatial prediction.

2.2.2 Wind Data

There are three sets of available wind data to be used in hurricane forecasting and hindcasting. Data provided by the National Hurricane Center (NHC) specifies the location of the storm center and the central pressure. The NHC provides forecasts every 6 hours on the location of the hurricane center (in degrees latitude/longitude) and sustained wind speeds (nautical miles) in different quadrants of the storm. For historical datasets such as these there is also "best" track information for the hurricane center and central pressure based on all available observations along the coast. The track information is necessary when creating wind fields based on the Holland wind model. Linear interpolation is used to interpolate the best track information to estimate the storm center and central pressure at ten–minute time intervals.

In addition, data on wind speed and direction at over 20 buoy locations along the Eastern coast is available from NOAA's National Data Buoy Center (NDBC). For this study we use the ten minute average wind speed (m/s) adjusted to a common height of 10 meters above sea level and ten minute average wind direction. Ten minute averaging times are considered to be representative of the timescales typically associated with oceanic response to surface stress (Houston et al., 1999). The wind speed and direction are then converted to u and v components.

Finally, gridded wind field data is available from NOAA's Hurricane Research Division (HRD). These wind fields are a combination of surface weather observations from ships, buoys, coastal platforms, surface aviation reports, reconnaissance aircraft data, and geostationary satellites which have been processed to conform to a common framework for height (10 m), exposure (marine or open terrain over land) and averaging period (maximum sustained 1 minute wind speed). The processing and quality control of the data use accepted methods from micrometeorology and wind engineering and provide the best available near real–time gridded hurricane wind analyses. However it is important to note that validation studies have shown that the HRD wind analyses are highly variable in accuracy, depending on the quality and quantity of the observations used, and on the appropriateness of the underlying assumptions used to manipulate the observations. Houston et al. (1999) report the gridded wind fields have been found to have estimated errors of up to 10%–20%. For this reason although buoy observations are used to create the HRD wind fields the observed winds from the NDBC buoy network are also used as a second data source to provide an estimate of the bias in the HRD analysis fields. For more information on the HRD data see Powell et al. (1996) and Powell and Houston (1996). HRD data is typically available at 3 hour intervals. HRD gridded fields are provided at a resolution of approximately 6km $\times$ 6km with the number of grid points on the order of 10^4.

2.2.3 Statistical Model for Data

A statistical framework is used to account for the difference in uncertainty associated with buoy and HRD analysis data. As described in the previous section the observed winds at the buoys are considered more reliable compared to the HRD analysis data which has been shown to have potentially large biases. This does not mean that no bias exists in the buoy data. For example wave distortion effects have been found to result in a negative bias when deriving wind stress from low–level anemonometers (Large et al., 1995). All of the buoy observations used in this study have anemometer heights that are at least five meters above mean sea level. In fact more than half of the sites are at heights greater than ten meters and so should not be influenced by surface waves. We cannot hope to accurately estimate the bias of both data sources but only the bias of the HRD data with respect to the buoy data. We take this approach since past studies suggest that the magnitude of the bias in the HRD data, which includes satellite and ship data, is larger than the potential bias in the buoys.

Let $\mathbf{V}^a(\mathbf{s}, t) = (v^a(\mathbf{s}, t), u^a(\mathbf{s}, t))^T$ be the u and v components provided by the HRD analysis at time t and location $\mathbf{s}$. Let $\mathbf{V}^b(\mathbf{s}, t) = (v^b(\mathbf{s}, t), u^b(\mathbf{s}, t))^T$ be wind data from the buoys time t and location $\mathbf{s}$. For now we consider time fixed and suppress the time index.

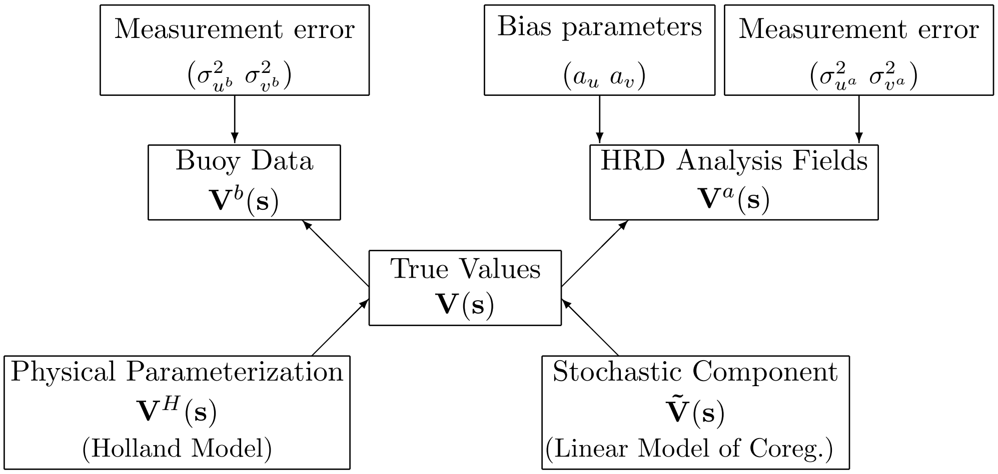

Figure 2.2: Model framework for wind components.

The observed data is modeled as a function of the underlying true (unobserved) wind process $\mathbf{V}(\mathbf{s}) = (u(\mathbf{s}), v(\mathbf{s}))^T$. Then $\mathbf{V}(\mathbf{s})$ is modeled as a multivariate spatial process with a mean function equal to the nonlinear parametric Holland formula. This framework is depicted in Figure 2.2. The data model includes bias parameters, $\mathbf{a}(\mathbf{s})$ and $\mathbf{b}(\mathbf{s})$, as well as measurement error terms, $\epsilon^b(\mathbf{s})$ and $\epsilon^a(\mathbf{s})$:

$$\hat{\mathbf{V}}^b(\mathbf{s}) = \mathbf{V}(\mathbf{s}) + \epsilon^b(\mathbf{s}) \tag{2.4}$$

$$\hat{\mathbf{V}}^a(\mathbf{s}) = \mathbf{a}(\mathbf{s}) + \mathbf{b}(\mathbf{s})\mathbf{V}(\mathbf{s}) + \epsilon^a(\mathbf{s}) \tag{2.5}$$

Here $\mathbf{a}(\mathbf{s}) = (a_u(\mathbf{s}), a_v(\mathbf{s}))^T$ and $\mathbf{b}(\mathbf{s}) = diag(b_u(\mathbf{s}), b_v(\mathbf{s}))$ are spatial–temporal functions for additive and multiplicative bias. These may be modeled as polynomial functions of location, spline functions or as a function of additional covariates such as temperature or pressure. Based on exploratory analysis comparing the buoy and HRD data we find that a multiplicative bias term in not necessary and set $b_u(\mathbf{s})$, $b_v(\mathbf{s}) = 1$ for all locations $\mathbf{s}$. Due to the limited amount of buoy data we consider only constant bias terms $a(\mathbf{s}) = (a_u, a_v)^T$ for all locations $\mathbf{s}$.

We assume both measurement error processes, ϵ^b and ϵ^a, are Gaussian white noise processes with $\epsilon^b(\mathbf{s}) \sim N(0, \Sigma^b)$, and $\epsilon^a(\mathbf{s}) \sim N(0, \Sigma^a)$ where $\Sigma^b = diag(\sigma^2_{u^b}, \sigma^2_{v^b})$ and $\Sigma^a = diag(\sigma^2_{u^a}, \sigma^2_{v^a})$. It is also assumed that these processes are independent of one another and the underlying process $\mathbf{V}(\mathbf{s})$. It is possible to instead model the measurement error as a spatially correlated process with some parametric spatial covariance function. We assume that the more complex spatial correlation structure is captured in the model for the underlying true process and hence there is no need for added complexity at this stage.

The statistical model for the unobserved "true" wind fields combines physical information from the parametric Holland formula with a multivariate spatial process. The Holland equations (2.1) – (2.3) define axis–symmetric wind fields, $\mathbf{V}^H = (u^H, v^H)^T$, under cyclostrophic balance and represent the large scale or mean structure of the surface winds. A multivariate spatial process, $\tilde{\mathbf{V}} = (\tilde{u}, \tilde{v})^T$, is used to model the variability of the observed data that cannot be captured by any parameterization of the Holland formula. The overall model for the underlying wind vector is defined as:

$$\mathbf{V}(\mathbf{s}) = \mathbf{V}^H(\mathbf{s}) + \tilde{\mathbf{V}}(\mathbf{s}) \tag{2.6}$$

More reliable wind fields can be obtained either by (1) introducing a more sophisticated deterministic wind model for $\mathbf{V}^H$ (e.g. Xie et al. (2006)) (2) developing a sophisticated statistical multivariate spatial model for the stochastic component $\tilde{\mathbf{V}}$. The proposed approach needs to be run in real–time so a balance between (1) and (2) is the ultimate solution to the forecasting problem. Alternatively, one could write a stochastic version of the deterministic model and approximate the physical model using a stochastic spatial basis. This is the approach of Wikle et al. (2001) for oceanic surface winds. Since the operational storm surge model uses parametric wind fields (essentially like the Holland winds), it is preferable to instead provide tools to better estimate the parameters of the Holland model accounting for uncertainty in the parameter estimates. This motivates the problem of defining a multivariate spatial model for the stochastic component. We adopt this approach here and evaluate the improvement in the wind prediction and storm surge forecast.

2.3 Modeling Multivariate Data

2.3.1 Linear Model of Coregionalization (LMC)

In our setting, any cross–dependency in the u and v winds is accounted for by using a multivariate spatial model based on a linear model of coregionalization. This approach was first introduced as a method of dimension reduction in order to model a set of p multivariate observations as a function of k unobservable factors with $k < p$. More recently, Schmidt and Gelfand (2003), Banerjee et al. (2004), and Gelfand et al. (2004) proposed using this approach to construct valid cross–covariance functions by modeling a set of dependent spatial processes as a linear combination of independent spatial processes. We briefly describe the LMC model followed by a description of the steps for Bayesian inference and prediction based on the available multivariate wind data.

We assume we are working with multivariate Gaussian processes. We find this is a reasonable assumption based on analysis of the observed u and v velocities after subtracting off the Holland mean function. In general, consider modeling p spatial processes $\tilde{\mathbf{V}}(\mathbf{s})_{(p\times 1)}$ for $\mathbf{s} \in \mathcal{D} \subset \mathcal{R}^2$. For a finite set of locations $\mathbf{s}_1, \ldots, \mathbf{s}_n$, let $\tilde{\mathbf{V}}$ denote a $np\times 1$ vector of the the p spatial functions evaluated at these n sites: $\{\tilde{\mathbf{V}} : \tilde{V}_1(\mathbf{s}_1), \ldots, \tilde{V}_1(\mathbf{s}_n), \ldots, \tilde{V}_p(\mathbf{s}_1), \ldots, \tilde{V}_p(\mathbf{s}_n)\}$. $\tilde{V}_j(\mathbf{s}_i)$ denotes the j^{th} spatial process of interest at location $\mathbf{s}$ with $j = 1, \ldots, p$ and $i = 1, \ldots, n$.

Let $w_j(\mathbf{s})$ be p latent independent Gaussian spatial processes with mean zero. Each process has covariance function $Cov(w_j(\mathbf{s}_i), w_j(\mathbf{s}_{i'})) = \rho_j(\mathbf{s}_i, \mathbf{s}_{i'}; \phi_j)$ with parameters ϕ_j, $j = 1, \ldots, p$. These covariance functions need not be stationary or isotropic. Now for $\mathbf{w}(\mathbf{s}) = (w_1(\mathbf{s}), \ldots, w_p(\mathbf{s}))^T$ the stochastic component in (2.6) is modeled as:

$$\tilde{\mathbf{V}}(\mathbf{s}) = \mathbf{A}\mathbf{w}(\mathbf{s}) \tag{2.7}$$

where $\mathbf{A}$ is a $p \times p$ weight matrix that determines the covariance between the p variables.

Without loss of generality assume $\mathbf{A}$ is a full rank lower triangular matrix. This implies $E(\tilde{\mathbf{V}}(\mathbf{s})) = \mathbf{0}$ and

$$Cov(\tilde{\mathbf{V}}(\mathbf{s}_i), \tilde{\mathbf{V}}(\mathbf{s}_{i'})) = \sum_{j=1}^{p} \rho_j(\mathbf{s}_i, \mathbf{s}_{i'}; \phi_j)\mathbf{T}_j \tag{2.8}$$

where $\mathbf{T}_j = \mathbf{a}_j\mathbf{a}_j^T$ and $\mathbf{a}_j$ is the j^{th} column of $\mathbf{A}$. Thus the cross–covariance is nonstationary, depending on the form of the ρ_j. Also the cross–covariance function is nonseparable, meaning we cannot separate the between site variability and the variability between measures taken at the same site.

$$Cov(\tilde{\mathbf{V}}, \tilde{\mathbf{V}}) = \sum_{j=1}^{p} \mathbf{T}_j \otimes \mathbf{R}_j \tag{2.9}$$

for $\mathbf{R}_j$ a set of p $n \times n$ matrices with elements $(R_j)_{ii'} = \rho_j(\mathbf{s}_i, \mathbf{s}_{i'}; \phi_j)$ and $\otimes$ denotes the Kronecker product. (Note, for $\tilde{\mathbf{V}} = (\tilde{V}(\mathbf{s}_1), \tilde{V}(\mathbf{s}_2), \ldots, \tilde{V}(\mathbf{s}_n))^T$ then the covariance is simply, $Cov(\tilde{\mathbf{V}}, \tilde{\mathbf{V}}) = \sum_{j=1}^{p} \mathbf{R}_j \otimes \mathbf{T}_j$.) For ρ_j stationary correlation functions, $\sum_{j=1}^{p} \mathbf{T}_j = \mathbf{T}$ is the covariance matrix for $\tilde{\mathbf{V}}(\mathbf{s})$ so there is a one–to–one relation between matrix $\mathbf{A}$ and matrix $\mathbf{T}$. For $p = 2$,

$$\mathbf{T} = \begin{pmatrix} a_{11}^2 & a_{11}a_{21} \\ a_{11}a_{21} & a_{21}^2 + a_{22}^2 \end{pmatrix} \tag{2.10}$$

Gelfand et al. (2004) also propose a nonstationary version of the LMC model by allowing the weight matrix to be a spatial function, $\mathbf{A}(\mathbf{s})$. Since there is a one–to–one correspondence between $\mathbf{A}(\mathbf{s})$ and the covariance matrix $\mathbf{T}(\mathbf{s})$, they develop an Spatial Wishart distribution for $\mathbf{T}(\mathbf{s})$. However this leads to a nonstandard distribution for $\tilde{\mathbf{V}}$ since the weight matrix in 2.7 is now a random spatial process. We thus propose to model all of the spatial structure within the latent processes. This will yield a nonstationary model for $\mathbf{V}(\mathbf{s})$ whose properties are easier to derive and interpret.

Note that a separable covariance model is just a special case of the LMC model where $\rho_j = \rho_{j'}$ for all $j, j' = 1, \ldots, p$ and ρ_j a stationary correlation function. In this case, $Cov(\tilde{\mathbf{V}}, \tilde{\mathbf{V}}) = \mathbf{T} \otimes \mathbf{R}$ for $(R)_{ii'} = \rho(\mathbf{s}_i, \mathbf{s}_{i'}; \phi)$. This is referred to as an intrinsic multivariate correlation model because the correlation structure of the set of variables is independent of the spatial correlation. That is, the correlation between two variables does not depend on spatial scale (Wackernagel, 1998):

$$Corr(\tilde{V}_j(\mathbf{s}_i), \tilde{V}_{j'}(\mathbf{s}_{i'})) = \frac{t_{jj'}\rho(\mathbf{s}_i, \mathbf{s}_{i'}; \phi)}{\sqrt{t_{jj}\rho(\mathbf{s}_i, \mathbf{s}_{i'}; \phi) t_{j'j'}\rho(\mathbf{s}_i, \mathbf{s}_{i'}; \phi)}} = \frac{t_{jj'}}{\sqrt{t_{jj}t_{j'j'}}} \tag{2.11}$$

In the following application we apply both a separable and nonseparable LMC model to the surface wind fields.

2.3.2 Statistical Methodology

We now turn to the issue of estimation and prediction of the wind fields to be used in the ocean model application. In order to estimate the wind field at a given set of locations we need estimates for the parameters $\mathbf{\Theta} = (\theta^D, \theta^H, \tilde{\theta})$ where θ^D=(a_u, a_v, $\sigma^2_{u^b}$, $\sigma^2_{v^b}$, $\sigma^2_{u^a}$, $\sigma^2_{v^b}$) are the bias and measurement error parameters, $\theta^H = (B, Rmax)$ are the parameters of the Holland model and $\tilde{\theta} = (A, \{\phi_j\})$ are the parameters for the LMC model for $j = 1, 2$. We present a Bayesian framework that allows for estimation of the parameters of the multivariate spatial model and the physically based parameterized wind function while accounting for potential additive bias in the observed data. Let $\hat{\mathbf{V}} = (\hat{\mathbf{V}}^b, \hat{\mathbf{V}}^a)$ be all the available wind data at time t from the buoys and HRD analysis wind fields and the Holland output at these sites. The posterior predictive distribution for the wind vector $(u(\mathbf{s}_0), v(\mathbf{s}_0))^T$ at a particular location $\mathbf{s}_0$ at time t is:

$$P(\mathbf{V}(\mathbf{s}_0)|\hat{\mathbf{V}}) \propto \int P(\mathbf{V}(\mathbf{s}_0)|\hat{\mathbf{V}}, \mathbf{\Theta}) \times P(\mathbf{\Theta}|\hat{\mathbf{V}}) d\mathbf{\Theta} \tag{2.12}$$

The full posterior predictive distribution (2.12) conditioned on the data from the n_a locations of the analysis wind fields and the n_b buoy sites can be approximated by using Markov Chain Monte Carlo. A Blocking Gibbs Sampling algorithm (within the software WinBUGS, Spiegelhalter et al. (1996)) is used to sample from the posterior distribution of the parameters by alternating between the parameters of the underlying process, $(\theta^H, \tilde{\theta})$, and the parameters for the measurement error and bias, θ^D. Our Gibbs algorithm proceeds in the following stages.

Stage 1:

Conditioned on the true wind process $\mathbf{V}$ we obtain the posterior distribution of the parameters, θ^D, that explain the bias and uncertainty about the data. The posterior distribution will be completely specified once priors are defined for θ^D:

$P(\hat{\mathbf{V}}^b|\mathbf{V}, \sigma^2_{u^b}, \sigma^2_{v^b}) \sim N(\mathbf{V}, \Sigma^b)$

$P(\hat{\mathbf{V}}^a|\mathbf{V}, a_u, a_v, \sigma^2_{u^a}, \sigma^2_{v^a}) \sim N((a_u, a_v)^T + \mathbf{V}, \Sigma^a)$

Stage 2:

At the second stage the statistical model for the underlying true winds is used. Let $\mathbf{V}^H$ denote the output from the Holland model with parameters θ^H at the locations of the observed data. Thus based on the conditional distribution:

$P(\mathbf{V}|\theta^H, \tilde{\theta}, \mathbf{V}^H) \sim N(\mathbf{V}^H, \sum_{j=1}^2 \mathbf{T}_j \otimes \mathbf{R}_j)$

and priors for $(\theta^H, \tilde{\theta})$, the posterior distribution is obtained for the parameters in the Holland model and for the parameters in the stochastic model.

Stage 3:

Values of $\mathbf{V}$, the unobserved true process, are simulated at the n_a locations of the analysis wind fields and the n_b buoy sites conditioned on values of $\mathbf{\Theta}$ updated in the previous two stages. Let $N = n_b + n_a$. The conditional multivariate distribution of $\mathbf{V}$ at a new location $\mathbf{s}_o$ is:

$$P(\mathbf{V}(\mathbf{s}_o)|\hat{\mathbf{V}}, \mathbf{\Theta}) \sim N(\mathbf{V}^H(\mathbf{s}_o) + \tau^T\Sigma^{-1}[\hat{\mathbf{V}} - \hat{\mu}], \hat{T} - \tau^T\Sigma^{-1}\tau) \tag{2.13}$$

where $\tau = \hat{Cov}(\mathbf{V}(\mathbf{s}_o), \hat{\mathbf{V}})$ is a $(2N\times2)$ matrix and $\Sigma_{(2N\times2N)} = \hat{Cov}(\hat{\mathbf{V}}, \hat{\mathbf{V}})$. $\hat{\mu} = (\mathbf{V}^H, \mathbf{a} + \mathbf{V}^H)^T$ is the Holland mean function evaluated at the buoy and HRD locations accounting for the bias in the HRD dataset. For more details on combining different sources of spatial data for Bayesian prediction see Fuentes and Raftery (2005).

Thus after iterating through the three stages there are M samples from the posterior distributions for the parameters, $\{\boldsymbol{\Theta}^{(m)}\}$, $m = 1, \ldots, M$. The Rao–Blackwellized estimator (Gelfand and Smith, 1990) for (2.12) is

$$P(\mathbf{V}(\mathbf{s}_o)|\hat{\mathbf{V}}) = \frac{1}{M} \sum_{m=1}^{M} P(\mathbf{V}(\mathbf{s}_o)|\hat{\mathbf{V}}, \boldsymbol{\Theta}^{(m)}) \tag{2.14}$$

Prior Specification

Uniform priors are used for the random effects such that $\sigma_{u^b}, \sigma_{v^b} \sim Uniform(.01, 1)$ and $\sigma_{u^a}, \sigma_{v^a} \sim Uniform(.01, 5)$. We use relatively informative priors for the buoy data here based on previous experience and analysis of similar datasets. These hyperprior values are based on information reported by the NDBC on the accuracy of the buoy equipment. The bias parameters are assigned normal priors with $a_u, a_v \sim N(0, .001)$ (.001 is the precision). Priors for the process parameters, $\theta^H = (B, Rmax)$, are chosen based on previous studies. $B \sim$ Uniform(1,2.5), $Rmax \sim$ Uniform(15,55) based on the physical constraints described earlier and discussed in Holland (1980). The next section describes the spatial covariance used for the hurricane case study. For our application, ϕ_1, ϕ_2 are range parameters which are assigned noninformative uniform priors, $Uniform(10, 1000)$, based on the size of the domain (in km) for the data.

Finally a prior is specified for the weight matrix $\mathbf{A}$. A conditional or hierarchical modeling approach described in Schmidt and Gelfand (2003) is used in order to ease the computational cost of working with the full $np \times np$ cross–covariance matrix of the stochastic component, $Cov(\tilde{\mathbf{V}}, \tilde{\mathbf{V}}) = \sum_{j=1}^{p} \mathbf{T}_j \otimes \mathbf{R}_j$, in Stage 2 of the algorithm. Evaluating (2.9) for $p = 2$, the joint covariance for $\tilde{\mathbf{V}}$ can be written as:

$$Cov(\tilde{\mathbf{V}}, \tilde{\mathbf{V}}) = \begin{pmatrix} a_{11}^2 \mathbf{R}_1 & a_{11} a_{21} \mathbf{R}_1 \\ a_{11} a_{21} \mathbf{R}_1 & a_{21}^2 \mathbf{R}_1 + a_{22}^2 \mathbf{R}_2 \end{pmatrix} \tag{2.15}$$

Now following a standard multivariate normal result (Appendix A):

$$\tilde{\mathbf{u}} \sim N(\mathbf{u}^H, a_{11}^2 \mathbf{R_1}) \tag{2.16}$$

$$\tilde{\mathbf{v}}|\tilde{\mathbf{u}} \sim N(\mathbf{v}^H + \frac{a_{21}}{a_{11}}(\tilde{\mathbf{u}} - \mathbf{u}^H), a_{22}^2 \mathbf{R}_2) \tag{2.17}$$

Using (2.10), (2.16)-(2.17) can be rewritten in terms of the elements of the $\mathbf{T}$ matrix. Under the new conditional parameterization the scale parameters $a_{11}^2 = t_{11}$ and $a_{22}^2 = t_{22} - \frac{t_{12}^2}{t_{11}}$ are assigned inverse gamma distributions ($\frac{1}{a_{11}^2}$, $\frac{1}{a_{22}^2} \sim Gamma(.01, .01)$) and the weight parameter $\frac{a_{21}}{a_{11}} = \frac{t_{12}}{t_{11}} \sim N(0, .001)$ (.001 is the precision). A similar approach can be used for multivariate data with $p > 2$. See the Schmidt and Gelfand (2003) study of multivariate air pollutant data for the conditional and unconditional models for $p = 3$.

2.4 Analysis of Data from Hurricane Charley

In this section we apply the multivariate statistical model for hurricane surface winds to a case study of Hurricane Charley of 2004. This case study was chosen because the hurricane surface winds induced measurable storm surge along the coast of Georgia and the Carolinas. We compare four different statistical models through a series of model diagnostics and look at posterior parameter estimates for the separable LMC model for 6 time periods. An empirical Bayesian estimation procedure is used for creating the gridded wind field inputs for the ocean model. Finally the numerical ocean model is run using the estimated wind fields and we compare the results to the predicted surge based on the Holland winds.

Water level gauges along the coast of Georgia and the Carolinas reported storm surge heights up to 1.5 meters above normal tidal levels. Gridded wind field data from the Hurricane Research Division is used for 6 times on August 14th (0200, 430, 730, 1030, 1330, and 1630 UTC) leading up to Hurricane Charley's landfall in South Carolina. Contour plots of the HRD wind speed (m/s) fields are shown in Figure 2.3. For computational reasons the gridded HRD data is sub–sampled to reduce the number of grid points to the order of 10^2. Buoy

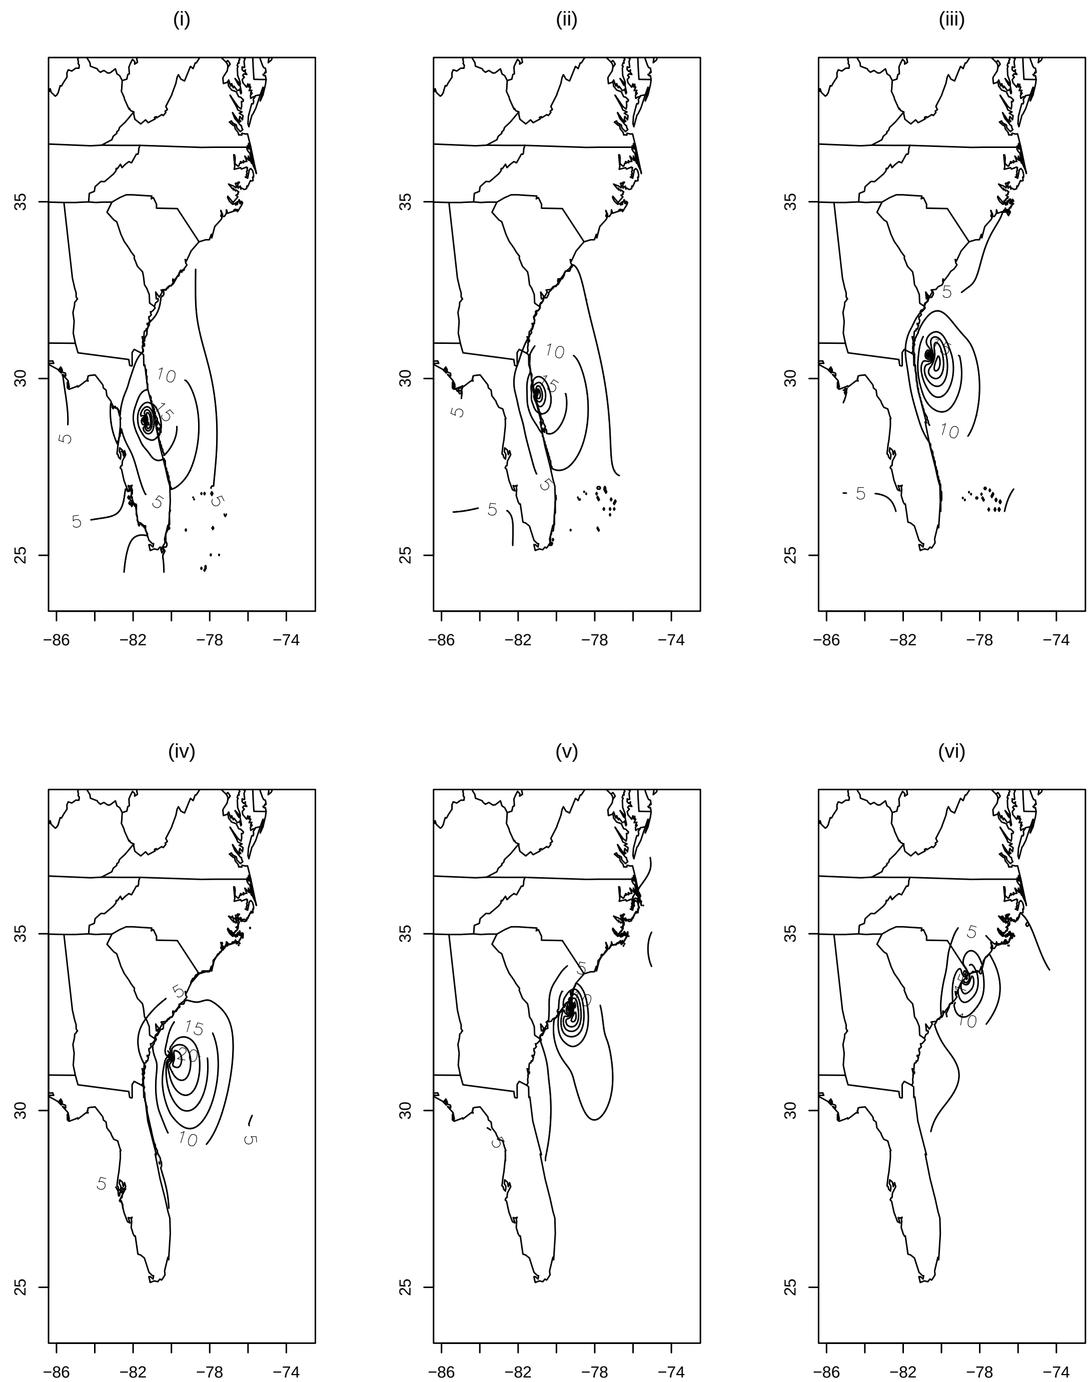

Figure 2.3: Wind contours (m/s) for August 14th hour (i) 200 UTC, (ii) 430 UTC, (iii) 730 UTC, (iv) 1030 UTC, (v) 1330 UTC, (vi) 1630 UTC.

observations from 23 locations are also used at these time periods to estimate potential bias in the HRD data.

2.4.1 Model Estimation for HRD and Buoy Data

Our objective is to fit a joint model for the wind field vectors and obtain the posterior predictive distribution of the winds at the ocean model grid points. Using the software GeoBugs (Spiegelhalter et al., 1996), four different models are fit using 225 locations of the HRD gridded data for u and v and data from the 23 buoys. Model 1 uses only the Holland mean function, i.e. $\tilde{\mathbf{V}}(\mathbf{s})$ is neglected. Note that in this case the parameters of the Holland winds are estimated using our framework rather than fixed values based on "expert knowledge". Model 2 treats the u and v components as independent so that $t_{12} = a_{11}a_{21} = 0$ in equation 2.15. Model 3 is a separable LMC model outlined in section 2.3.1 so that the parameters of the underlying spatial processes are assumed to be equal (i.e. $\phi_1 = \phi_2$). Finally, Model 4 is a nonseparable LMC model which does not impose this assumption about the spatial processes.

For this application, $\rho_j(\mathbf{s}_i - \mathbf{s}_{i'}; \phi_j)$ are assumed to be stationary correlation functions with parameters ϕ_j; $j = 1, 2$. We did allow for geometric anisotropy, e.g. Fuentes et al. (2005), but found no significant evidence of a lack of isotropy in the u and v winds for this case. Here we use an exponential correlation function such that $\rho_j(d; \phi_j) = exp(-d/\phi_j)$, $j = 1, 2$ for $d = ||\mathbf{s}_i - \mathbf{s}_{i'}||$ and the range parameter ϕ_j represents the distance at which two wind measurements become practically uncorrelated. Note that ϕ_1, the range parameter for the underlying spatial process $w_1(\mathbf{s})$, is also the range for the u wind component. We estimate the true range with the effective range parameter $\hat{\phi}_u$ that solves $\rho_1(\phi) = .05$. However to find the range parameter for the v component, the correlation function is the weighted average correlation (Gelfand et al., 2004):

$$\rho_v(h) = \frac{a_{21}^2 \rho_1(h) + a_{22}^2 \rho_2(h)}{a_{21}^2 + a_{22}^2} \tag{2.18}$$

Thus given the parameters $\mathbf{A}$, ϕ_1 and ϕ_2 we can find the effective range $\hat{\phi}_v$, by using an iterative root solving algorithm to solve $\rho_v(\phi) = .05$.

Table 2.1: Posterior medians and 2.5 and 97.5 percentiles for the Holland function and bias parameters in the separable LMC spatial model for u and v.

		B			Rmax	
UTC	2.5%	Median	97.5%	2.5%	Median	97.5%
200	1.62	2.29	2.50	5.00	5.10	5.90
430	1.04	1.35	1.83	5.86	7.11	8.69
730	1.27	1.65	2.22	20.65	23.39	28.43
1030	1.00	1.35	2.45	6.65	12.96	27.44
1330	1.00	1.13	1.60	7.46	10.84	27.45
1630	1.46	1.91	2.45	5.41	7.25	10.87

		a_u			a_v	
UTC	2.5%	Median	97.5%	2.5%	Median	97.5%
200	-.56	-.12	.33	-.83	-.31	.19
430	.38	.76	1.14	-.86	-.02	.35
730	-.07	.44	.89	-1.58	-.95	-.28
1030	-.28	.27	.73	1.45	1.85	2.32
1330	.92	1.41	1.86	1.19	1.62	1.99
1630	-.33	.29	.89	-1.3	-.91	-.45

Table 2.1 shows the posterior distribution summaries of the Holland parameters in the separable LMC model for u and v. Here we are using a temporally varying spatial model by allowing the parameters of the spatial model to change over time. That is $\Theta_t = (\theta_t^H, \tilde{\theta}_t)$. The parameter estimates are significantly different for different time periods. Standard values for the Holland model are B=1.9 and $Rmax$=35 to 55 km depending on the size and intensity of the storm, e.g. see Hsu and Yan (1998). This range of values for $Rmax$ based on previous

studies does not capture any of the values estimated in this case study. As shown in Figure 2.3, at 200 and 430 UTC the hurricane has just passed over Florida and is less symmetric. The radius of maximum winds ($Rmax$) at these times is 5 to 7 km. Using the standard values greatly overestimates the strength and extent of the winds at such times when the storm is less organized.

Table 2.1 also gives the posterior summaries of the bias parameters, a_u and a_v. In general the bias in the HRD winds relative to the buoy observations tends to be less than 1 m/s. The largest bias is seen in the North–South winds at hours 1030 and 1330 UTC. Another parameter of interest is the correlation between the wind components. The posterior estimates of the cross–correlation parameter ($t_{12}/\sqrt{t_{11}t_{22}}$) range from .24 to -.23 across the six time periods. The posterior intervals for the correlation do not contain zero, suggesting there is significant correlation between the u and v winds at all time periods.

Model Diagnostics and Calibration

The Deviance Information Criterion (DIC) (Spiegelhalter et al., 2002) shown in Table 2.2 is used to evaluate the different statistical models (where lower DIC values are better). All of the spatial models perform much better than the model with just the Holland mean function. The multivariate LMC models fit the data better for all of the times except 430 and 1630 UTC. At 430 UTC the DIC values for the independent and nonseparable models are similar. At 1630 UTC the independent model appears to be the better fit. At 430 and 1630 UTC the eye of the hurricane is directly off the coast and the wind speed contours appear more symmetric than the other time periods. Thus although it is clear that the spatial models do a much better job of estimating the wind fields it appears that the issue of independence of the u and v velocity components depends on the structure and organization of the storm at a given time.

A second criterion for model choice is calculated for comparison based on the method

Table 2.2: Deviance Information Criterion (DIC) values for Hurricane Charley data.

Model	200 UTC	430 UTC	730 UTC	1030 UTC	1330 UTC	1630 UTC
Model 1	3151.8	3400.9	3433.7	3541.7	3147.8	3353.7
Model 2	2827.6	1402.9	1419.4	1551.7	1522.5	1939.2
Model 3	2093.6	1505.9	1292.4	1023.7	1381.2	2167.1
Model 4	1832.5	1417.3	1093.9	1285.0	1439.4	2547.1

proposed by Gelfand and Ghosh (1998). These values (based on the squared error loss function) are shown in Table 2.3. Once again there is a clear difference between Model 1 and the three spatial models. Overall the three spatial models have very similar values with Model 3 having slightly lower values except at 730 UTC and 1630 UTC.

Table 2.3: Model choice results for each time period using Models 1, 2, 3 and 4.

Model	200 UTC	430 UTC	730 UTC	1030 UTC	1330 UTC	1630 UTC
Model 1	19.13	19.48	19.93	25.79	18.61	19.73
Model 2	1.185	.6152	.5240	.5775	.5788	1.188
Model 3	.8988	.5505	.5914	.5447	.5000	1.346
Model 4	1.071	.5770	.6599	.5457	.5679	1.226

Cross validation is used to evaluate how well the models predict winds at a new location. Models 1 through 4 are used to predict at thirty HRD grid points that were not used for parameter estimation. The median values of the posterior predictive distributions are used to compare to the observed values at the hold out sites. Analysis of the cross–validation errors under each model shows no significant prediction bias for any of the models. The non–spatial model using only the Holland mean function has much larger error variance and several outlier values show that the winds at a few locations are underpredicted by up to 15 to 20 m/s. The root means square prediction error (RMSPE) values for Models 2, 3 and 4 (the spatial models) are very similar. The calibration is compared for each model and each time period based on the 95% posterior predictive intervals. For the models to be well calibrated the observed velocities should fall within the posterior predictive intervals at least 95% of the time. Figure 2.4 shows calibration plots for the v component for the hold out sites at 430 UTC. The prediction results for Models 2, 3 and 4 are practically identical. Similar results were found for the other time periods and for the u winds. All of the spatial models have calibration percents greater than or equal to 95% at all times. However the calibration percents for Model 1 (the non–spatial model) are lower, ranging from 86.7% to 96.7%. Thus although the BIC values show a clear difference in estimation results across the four models, it appears that Models 2, 3 and 4 perform well in terms of prediction. In fact it is not surprising that although the three spatial models result in different covariance parameter estimates this does not have a dramatic impact on spatial prediction (Stein, 1999).

2.4.2 Ocean Model Application

We now consider the application of the proposed wind field model for storm surge prediction using the Princeton Ocean Model (POM) described in Chapter 1. In the previous section we used the fully Bayesian solution for the posterior predictive distribution $P(\mathbf{V}(\mathbf{s}_o)|\hat{\mathbf{V}})$ outlined in section 2.3.2 to predict at 30 cross validation sites for six different time periods. However, for the storm surge application we must create predicted wind fields for every grid point of the ocean model domain for 96 ten–minute time increments over a period from 200 UTC to 1750 UTC on August 14, 2004. For this application the ocean model is typically run at a resolution of 2 to 4 km grids. In our case the model is run using a resolution of 1 minute longitude by 1 minute latitude grid ($\approx$ 1.5 km by 1.9 km) with a total of 301$\times$229 grid points. Hence it is computationally infeasible to use the fully Bayesian prediction to predict the winds at a such a high resolution.

Here we take an empirical Bayesian approach and use the posterior median values of the parameters, $\mathbf{\Theta}$, as the fixed "true" parameter values. Then the predictions are made

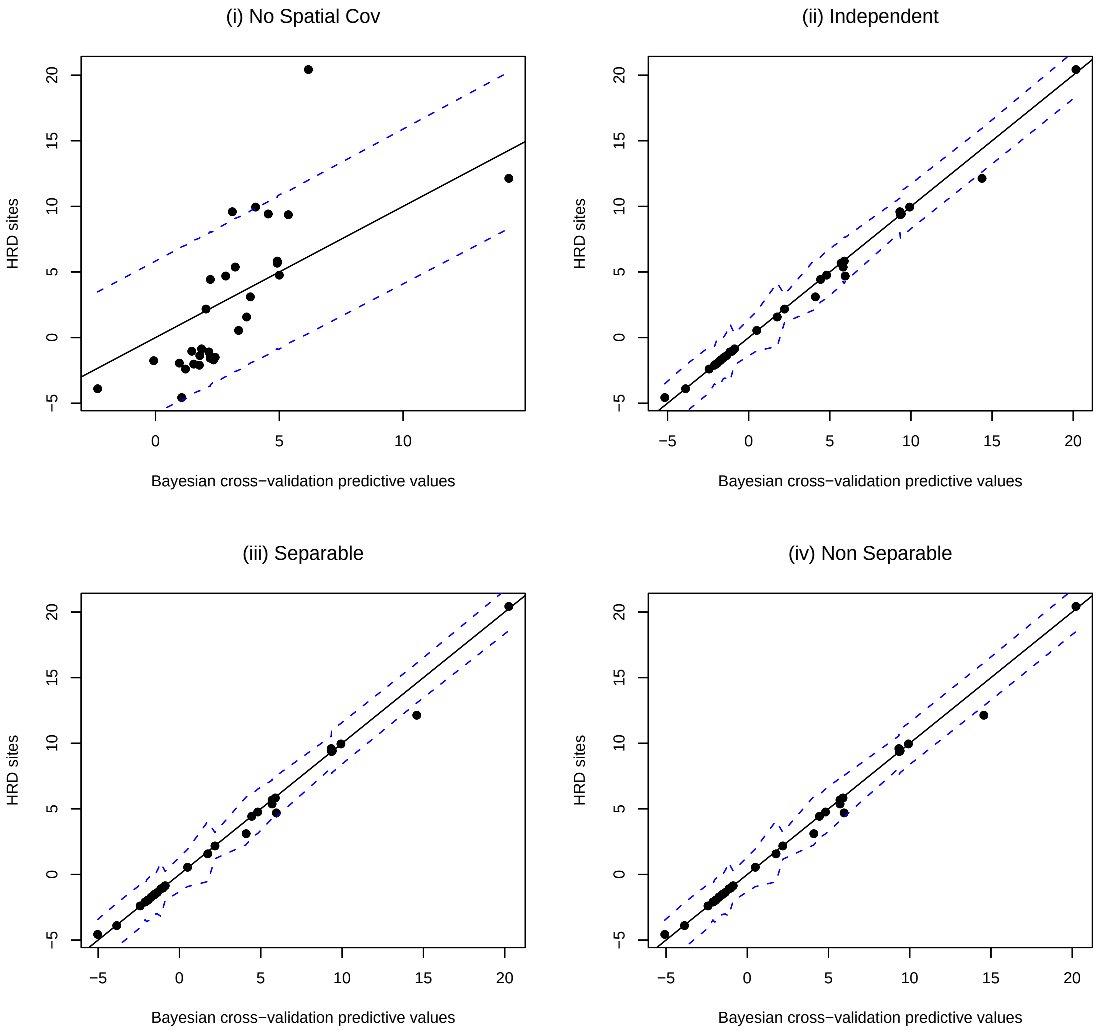

Figure 2.4: Cross Validation 430 UTC: Calibration plots for the v winds at 30 cross validation locations for 4 models: (i) Model 1 – Holland mean with no spatial covariance (ii) Model 2 – independent (iii) Model 3 – separable LMC (iv) Model 4 – nonseparable LMC.

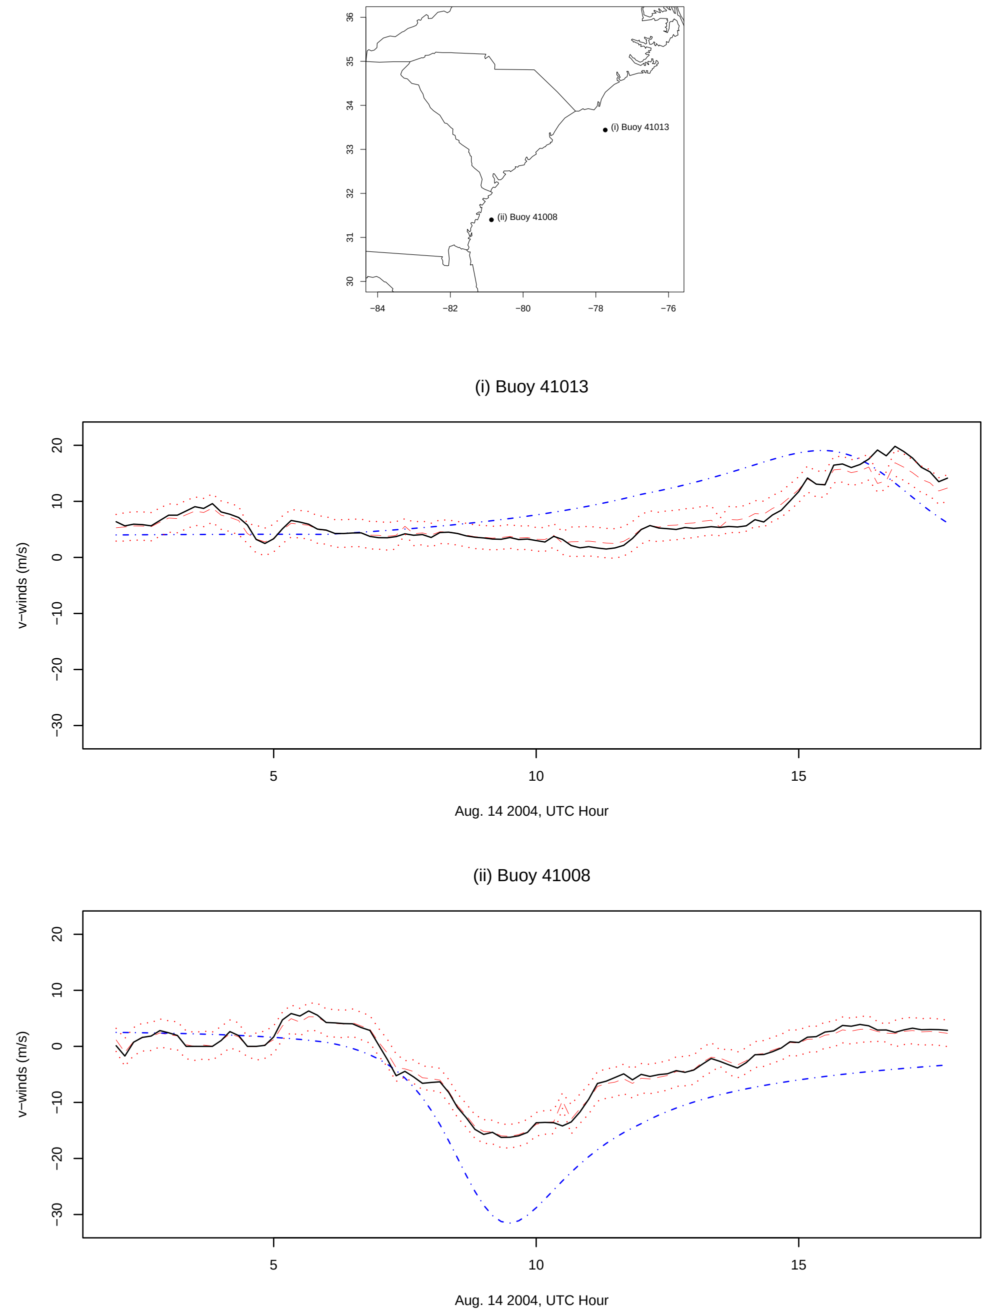

Figure 2.5: Map of 2 NBDC buoys within the POM domain. The remaining plots show the observed v winds in m/s at each buoy in solid black for August 14th 200 UTC to 1750 UTC. The dashed and dotted line is the Holland output based on fixed parameter values ($B = 1.9$, $Rmax = 46$). The dashed lines are the predicted values and the dotted lines are 95% posterior predictive intervals using Model 3.

following the conditional distribution in equation (2.13). The two methods are compared based on the 30 cross validation sites used in the previous section. Results show that the coverage is almost identical for both the fully Bayes and empirical Bayes prediction. Also, on average, the widths of the fully Bayesian credible intervals are only 4% larger than the intervals based on the empirical approach. Thus it appears that in this case the empirical Bayes prediction is a reasonable approach as well as being significantly more computationally efficient.

The separable LMC model is used for prediction although, as noted earlier, the other spatial models had almost identical results in terms of prediction. At the 10 minute intervals when buoy observations are available but not the HRD datasets the parameter values used for prediction are based on the estimation results from the last available HRD time. Let $t = 1, \ldots, m$ be the 10 minute intervals from August 14 200 UTC to August 14 1730 UTC, 2004. Let $\{t_H\}$, $H = 1, \ldots, 6$ be a subset of these discrete times representing the times when HRD data is available. Then $\theta^H(t) = \theta^H(t_H)$ and $\tilde{\Theta}(t) = \tilde{\Theta}(t_H)$ for $t_H \leq t$. In this way we have estimates at 96 different time periods corresponding to 10 minute intervals between 200 to 1750 UTC on August 14th. Figure 2.5 shows the observed and predicted v velocities for two buoys located within the ocean model domain. (Similar results were found for the u component.) The dashed and dotted line is the output from the Holland model using fixed parameter values ($B = 1.9$, $Rmax = 46$). The Holland function is unrealistically smooth and tends to exaggerate the change in the winds as the hurricane approaches these locations. There is also no measure of uncertainty in the values from the Holland function. The posterior predicted values (dashed lines) match closely to the observed data at each location. The corresponding 95% posterior predictive intervals are also included as dotted lines but we note that since an empirical Bayes approach is used here, these intervals do not account for the uncertainty in the model parameters. Still, the observed wind values fall outside of these intervals fewer than 6 of the 96 time periods in each plot.

Model Diagnostics

As noted before, the cross validation is repetaed at the 30 hold out sites using the empirical Bayesian prediction. The calibration percents were all greater than 95%. Table 2.4 shows the RMSPE values in meters per second for each time period. We compare the Holland model output with fixed parameters ($B = 1.9$, $Rmax = 46$), the non–spatial model and the separable LMC model using both the fully Bayesian and empirical Bayesian approaches. The Bayesian estimation of the Holland parameters decreases the RMSPE values by more than forty percent in most cases. The predictions based on the spatial model have much lower RMSPE values than the non–spatial model for all time periods and the values are very similar for both methods.

Table 2.4: RMSPE for $\mathbf{V}$ (m/s) based on 30 cross validation sites using the Holland model output (fixed parameters), fully Bayesian prediction and the empirical Bayesian approach based on the non–spatial model (Model 1) and separable LMC model (Model 3).

Model	200	430	730	1030	1330	1630
Holland Output	9.57	7.70	6.33	7.21	6.53	6.40
Full Bayes (Non–spatial)	3.62	3.30	3.51	4.55	3.66	3.05
Full Bayes (Spatial)	.54	.25	.84	1.79	.64	.18
Emp. Bayes (Spatial)	.51	.43	1.13	1.84	.62	.17

The prediction of the v winds at the two buoys for the 96 time periods is repeated, this time without using the observed winds at these two sites. Figure 2.6 shows the predicted values under Model 3 (plots i-a and ii-a) and Model 1 (plots i-b and ii-b). The spatial model is still able to better capture the variability in the winds when compared to the Holland model output. The observed v winds fall outside the 95% posterior prediction intervals for plots i-a and ii-a 26% and 7% of the time periods, respectively. The non–spatial model also shows an improvement over the Holland model output but the prediction intervals do not capture the observed values as often as the intervals based on the spatial model. Thus there is a clear

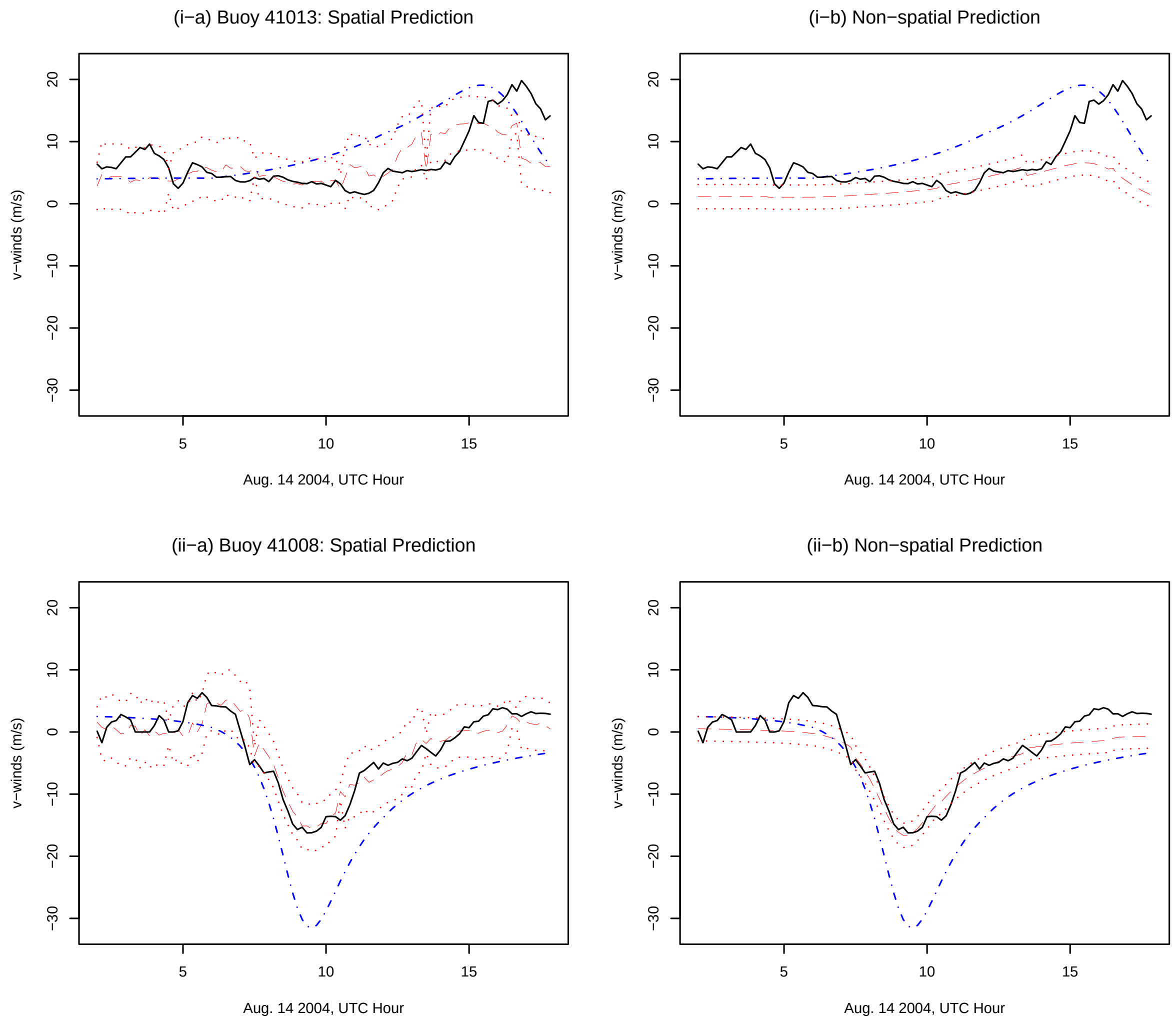

Figure 2.6: Cross validation time series plots: Observed v winds in m/s at each buoy in solid black for August 14th 200 UTC to 1750 UTC. The dashed and dotted line is the Holland output based on fixed parameter values ($B = 1.9$, $Rmax = 46$). The dashed line is the predicted values and the dotted lines are 95% credible intervals under (i-a), (ii-a) Model 3 and (i-b), (ii-b) Model 1.

advantage in using the statistical framework to estimate the Holland model parameters and the uncertainty in these parameters. Furthermore, there is added value in incorporating a multivariate spatial covariance through the stochastic component.

2.4.3 Storm Surge Prediction

The predicted wind fields are used as the input fields for the POM model to spin up and force the ocean model for a time between August 14th hour 0 UTC and hour 1800 UTC. Observed water levels were used from a series of coastal water level stations maintained by NOAA's National Ocean Service (NOS) Center for Operational Oceanographic Products and Services (CO–OPS). We have hourly data at 7 coastal sites for Hurricane Charley for the surface water elevation in meters shown in Figure 2.7. These observations are adjusted to account for differences in the tidal gauge locations and the effect of daily tidal levels to determine the change in water elevation that can be attributed to storm surge. Figure 2.7 shows the observed water levels at each of these sites as Hurricane Charley moves along the coast. The Sunset Beach site shows the greatest storm surge of 1.67 meters.

The dashed and dotted line shows the predicted water elevation at each site based on hourly output of the Princeton Ocean Model under the original Holland wind forcings. Using the Holland winds with fixed parameter values the ocean model tends to greatly overestimate the storm surge at almost every locations. The dotted line shows the predicted values using the statistically based wind estimates, $\mathbf{V} = \mathbf{V}^H + \tilde{\mathbf{V}}$, where $\tilde{\mathbf{V}}$ is based on Model 3 (separable LMC). The prediction error is reduced and at several of the sites these estimates tend to better capture the timing and magnitude of the peak storm surge. However the model output based on the LMC winds still underestimates the greatest surge seen at the Sunset Beach site at 1600 UTC.

2.4.4 Comparison of Data Assimilaiton Approaches

In this section the different data assimilation methods outlined in Chapters 1 and 2 are compared. The prediction errors from the EnsKF are compared to the errors from the ocean model output using the improved wind field inputs. Figure 2.8 shows a series of boxplots corresponding to the prediction errors from the control run and using the different data

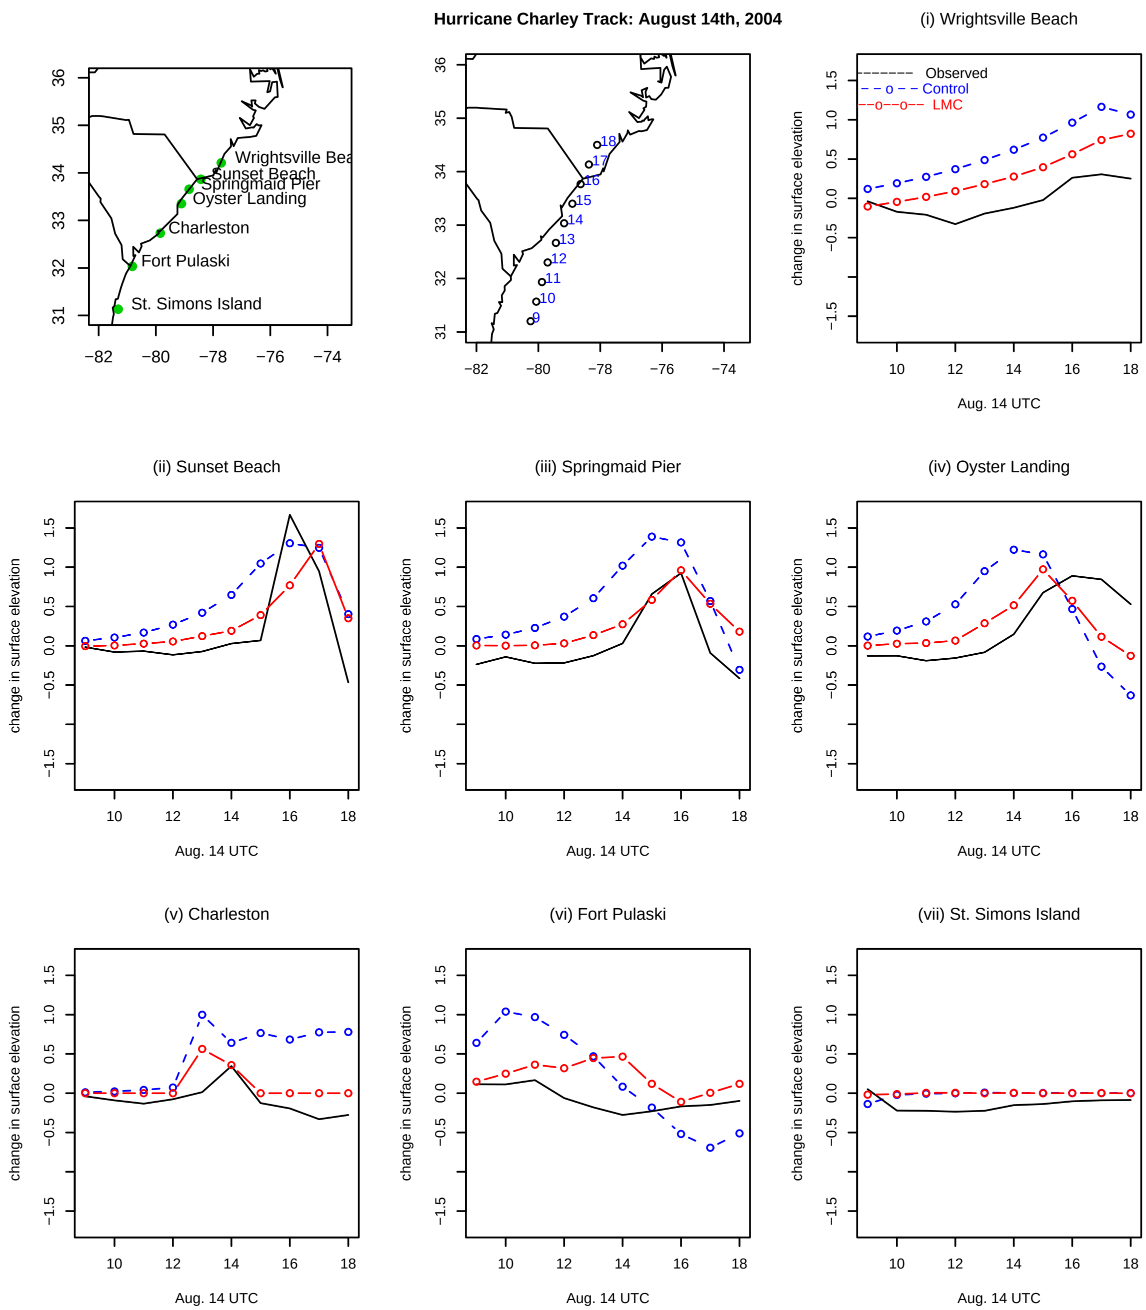

Figure 2.7: Maps show the location of 7 coastal water elevation gauges and the track of Hurricane Charley on August 14th for hours 900 UTC to 1800 UTC. Time series plots show observed water elevation at each site as a solid black line. The dotted line shows the predicted elevation using the predicted winds from the LMC statistical model to force the ocean model run. The dashed and dotted line shows the estimates when the ocean model is forced using the original Holland wind model.

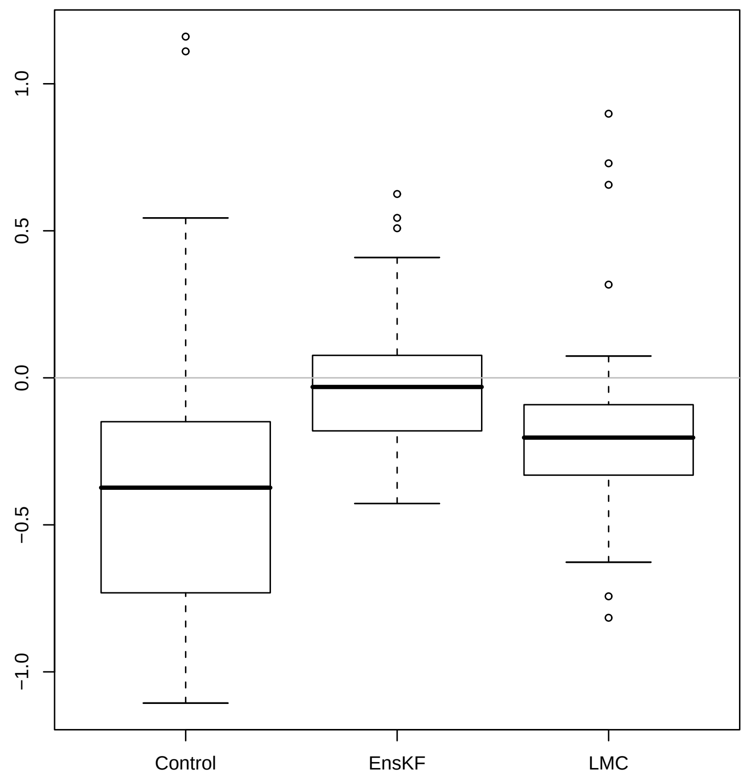

Figure 2.8: Boxplots of residuals (observed elevation minus predicted elevation) for control runs, EnsKF results and output based on LMC winds. Plots are based on 70 observations using all ten hourly values at the seven coastal locations.

assimilation methods. These plots are based on seventy observations using the ten hourly elevation values at the seven coastal locations. Both assimilation methods show a decrease in error bias and in error variability when compared to the control run. The EnsKF has a lower bias and fewer large errors.

Reviewing the time series plots in Figures 1.3 and 2.7, we can make several observations. The surge at the Oyster Landing site is predicted best by the EnsKF using the inundation model which is able to model the prolonged flooding in this area as the hurricane moves inland. The improved wind inputs also improve the prediction of storm surge and provide better estimates of the magnitude of the surge compared to the control run. The control run and both estimates miss the maximum surge at the Sunset Beach site.

2.5 Discussion

In this chapter we present a flexible hierarchical framework to incorporate observations from different data sources with physical knowledge of hurricane surface wind fields. The multivariate spatial model proposed here allows for cross dependency between the two wind components. We found that including this dependency was important for estimation at the time periods when the hurricane structure was more asymmetric. We used empirical Bayes prediction for computational reasons but clearly a fully Bayesian analysis provides slightly better coverage of the credible intervals. The cross validation analysis shows that prediction results are similar under either approach.

All of the spatial models assumed that the underlying univariate processes of the stochastic residual component were Gaussian and stationary. Based on exploratory analysis we found that the Gaussian assumption was reasonable. In terms of prediction the three spatial models, Model 2, Model 3, Model 4, produced similar results and were able to produce very accurate predictions when compared to the observed data. We also saw a clear improvement in prediction when the Holland parameters were estimated using the statistical framework proposed here (Model 1) when compared to the results based on "expert knowledge".

Another approach to capture more complex spatial structures would be to allow for potential nonstationarity in the spatial processes. As noted earlier, the proposed methodology can still be used in this case by allowing for the underlying univariate processes to be nonstationary. Also of interest would be to consider a full spatial–temporal covariance model rather than the temporally evolving spatial model used here. In the following chapter we propose a new spatial–temporal LMC model that can handle both of these extensions.

In applications such as storm surge prediction it is often unclear whether prediction errors are a result of simplified or neglected physical processes used in the numerical model, inaccuracies in the initial conditions or errors in the forcing fields. As was investigated here we found that prediction errors in the storm surge estimates for Hurricane Charley are clearly a result of a combination of these factors. Improving the wind fields inputs was shown to improve the estimated water levels at seven coastal sites as was using the EnsKF with a higher resolution version of POM coupled with an inundation model.

Coastal ocean prediction for hurricane applications remains a challenging problem since we must balance potential improvement in estimation and prediction from a more sophisticated model and computational cost. In terms of computational cost, the greatest cost of the methods presented for the surface wind fields is from the Gibbs sampling algorithm required for the estimation of the parameters of the statistical model whenever HRD data is available. This step takes approximately fifty minutes for 1000 updates using the WinBUGS software. However the prediction step using the empirical Bayes approach takes only minutes to create the gridded wind inputs based on all available buoy data at ten minute intervals. Thus after an initial lag when new HRD data becomes available it is possible to run the ocean model with updated winds for real–time storm surge prediction. In general the EnsKF approach is computationally more efficient. Using the high resolution inundation ocean model it takes aproximately eighteen minutes to forecast a twenty member ensemble forward one hour and then calculate the EnsKF update using ten machines on the PAMS College Linux cluster.

Chapter 3

A New Multivarite Nonstationary Spatial Temporal Model

3.1 Introduction

Statistical space–time modeling has proven an essential tool in the environmental sciences to describe complex spatial and temporal behavior of physical processes. Many classes of space–time models rely on the assumption that the processes of interest are stationary, meaning that the structure of the spatial–temporal variability does not change over the spatial–temporal domain. Another common assumption is that there is no interaction in the spatial and temporal variability, i.e. the space–time covariance is separable. These assumptions have been shown to be unrealistic for numerous applications such as modeling coastal wind velocities, precipitation amounts or air pollutant concentrations. As a result new classes of statistical models have been developed to address these issues.

In this chapter we discuss the issue of space–time modeling of multivariate observations. We review the methods for specifying a cross–covariance function that also account for potential nonstationarity in the spatial data. We propose a new class of nonstationary multivariate models based on an extension of the linear model of coregionalization (LMC). This model provides a flexible tool for predicting in space and time. In addition the statistical model has the advantage that the estimated parameters are easy to interpret from a geostatistical point of view. Such estimation results are relevant in applications where there is interest in better understanding how different physical processes interact and how these interactions change over time and space.

In our application we jointly model the orthogonal wind components of hurricane surface winds. When modeling surface winds the assumption of stationary covariance may be unrealistic. In the previous chapter we describe how it is known hurricane winds are not symmetric. In some cases this may imply the need for a nonstationary covariance model. We estimate the parameters of a deterministic mean function for these wind fields while using a cross–covariance model to understand how the orthogonal wind components are related and how these processes have different spatial variability. A nonstationary multivariate space–time model is used to analyze surface wind field data of hurricane Floyd from September of 1999. We relate the parameter estimates to physical features of the hurricane. The nonstationary model is also shown to perform well in terms of prediction. Wind field data is provided by NOAA's Hurricane Research Division (HRD) and is used to improve hurricane forecasts for the Eastern United States coastline. Buoy data is also available at several locations along the coast. While the HRD data is spatially dense it is only available every three hours. The buoy data provides wind observations every ten minutes. Thus a Bayesian spatial–temporal framework allows us to utilize both sources of information for improved prediction.

Various nonstationary spatial approaches have been applied to univariate processes. Gelfand et al. (2004) review the multivariate versions of several stationary and nonstationary methods including kernel convolution and convolution of covariance functions. They propose a new nonstationary approach called the spatially varying LMC (SVLMC). We discuss the extensions of several of these approaches in further detail in section 3.2. De Iaco and D. E. Myers (2003) propose a multivariate space–time model based on a nonseparable space–time variogram model. The spatial variability is assumed to be stationary and the proposed model is estimated using the sample marginal space and time variograms and graphical approaches. We describe a fully Bayesian framework for fitting the nonstationary multivariate model presented here.

Le and Zidek (1992) propose a hierarchical Bayesian approach for spatial–temporal interpolation by treating the cross–covariance matrix as a random realization from an inverse Wishart distribution. Their approach is further developed by Brown et al. (1994), Le et al. (1997), Sun et al. (1998) and Kibria et al. (2002) for interpolating multivariate air pollutant values at ungauged sites. As noted by Gelfand et al. (2004), the resulting covariance matrix is nonstationary and positive definite but no longer associated with a spatial process since the elements of the covariance matrix are no longer a function of the spatial locations. Since our goal is to provide an intuitive parameterization of the spatial–temporal cross–covariance we rely on covariance functions that are valid by construction and are associated with traditional univariate spatial and temporal covariance models.

Other spatial–temporal models have been applied to surface wind field data. For example Fuentes et al. (2005) found that surface wind fields near the Chesapeake bay area exhibited different patterns of spatial dependency over land when compared to the winds over the water. Wikle et al. (2001) present a hierarchical spatio–temporal model for surface wind fields over the tropical oceans. They use a thin–fluid approximation for large–scale tropical dynamics to represent the evolution of oceanic surface wind fields over time. Here we do not have such a dynamic physical (deterministic) model to represent the evolution of hurricane winds over time. The Holland formula incorporates temporal information in that it is a function of the storms central pressure and location which will change over time. Thus we use the statistical modeling approach outlined in Chapter 2 with the additional flexibility of a spatial–temporal covariance applied to the stochastic component of the underlying wind vector (i.e. $\tilde{\mathbf{V}}$ in (2.6)).

The model presented here allows for cross dependency between the wind components as well as potential nonstationarity or nonseparability in the specification of the spatial–temporal covariance. Section 3.2 reviews nonstationary multivariate models based on kernel convolution and LMC methods. Section 3.3 describes a new nonstationary multivariate model and a Bayesian approach to estimation and prediction. Section 3.4 is an application to hurricane surface wind field data from hurricane Floyd. We show that the multivariate nonstationary model is able to better capture the spatial variability of the wind fields in different quadrants of the storm. A discussion is given in the concluding section 3.5.

3.2 Modeling Nonstationary Multivariate Data

Rather than assuming independence between observations of different variables we implicitly account for any dependence through the specification of a cross–covariance function. In general let $\mathbf{V}(\mathbf{s}, t)$ be a $p \times 1$ vector valued spatial–temporal process. For r locations $\mathbf{s}_1, \ldots, \mathbf{s}_r \in \mathcal{D}$ and m observations over time $t_1, \ldots, t_m \in \mathcal{T}$ we have $n = r \times m$ observations for each variable and np total observations. Let the $(np \times 1)$ data vector be defined such that $\mathbf{V} = (\mathbf{V}_1^T, \ldots, \mathbf{V}_p^T)^T$ where $\mathbf{V}_j = (V_j(\mathbf{s}_1, t_1), V_j(\mathbf{s}_1, t_2), \ldots, V_j(\mathbf{s}_1, t_m), V_j(\mathbf{s}_2, t_1), \ldots, V_j(\mathbf{s}_r, t_m))^T$. The most basic cross–covariance model is separable in space and time and also separable in the sense that the correlation structure between different variables does not depend on spatial scale or temporal scale. Thus the $np \times np$ cross–covariance matrix for $\mathbf{V}$, $\mathbf{\Sigma}$, can be written as a Kronecker product ($\otimes$) of the $p \times p$ variance/covariance matrix of the p variables, the $r \times r$ spatial covariance matrix and the $m \times m$ temporal covariance matrix.

$$\mathbf{\Sigma} = \Sigma_V \otimes \Sigma_S \otimes \Sigma_T \tag{3.1}$$

Although this model simplifies computation for estimation and prediction it is limited by the fact that it does not allow for spatial–temporal interactions or for the correlation between variables to evolve over space or time. In section 3.3 we propose a space–time cross–covariance model based on modeling a set of p dependent spatial–temporal processes as a linear combination of p independent processes. We do not assume stationarity and separability in the spatial–temporal cross–covariance but show that these are special cases of the general model. In the following sections we review the methods proposed by Higdon et al. (1998) and Fuentes and Smith (2001) for nonstationary spatial modeling followed by a description of the SVLMC approach proposed by Gelfand et al. (2004).

3.2.1 Spatially Varying Kernel Convolution

Convolution methods are a standard approach for modeling univariate stationary Gaussian spataial processes (Yanglom, 1962). Consider a white noise process $z(\mathbf{s})$ defined for any regions A, $B \in \mathcal{R}^2$ such that

$$\begin{aligned} Z_A &= \int_A z(u)du \sim N(0, \sigma^2 area(A)) \\ Cov(Z_A, Z_B) &= \sigma^2 area(A \cap B) \end{aligned} \tag{3.2}$$

Any stationary Gaussian process $V(\mathbf{s})$ can be expressed as the convolution of $z(\mathbf{s})$ with kernel $k(\mathbf{s})$

$$V(\mathbf{s}) = \int k(\mathbf{s} - \mathbf{u}) z(\mathbf{u}) d\mathbf{u} \tag{3.3}$$

It can immediately be shown $V(\mathbf{s})$ has $E(V(\mathbf{s})) = 0$ and $Cov(V(\mathbf{s}_i), V(\mathbf{s}_{i'})) = \sigma^2 \int k(\mathbf{s}_i - \mathbf{u}) k(\mathbf{s}_{i'} - \mathbf{u}) d\mathbf{u}$. Using a change of variables we see $V(\mathbf{s})$ is stationary:

$$Cov(V(\mathbf{s}_i), V(\mathbf{s}_{i'})) = \sigma^2 \int k(\mathbf{d} + \mathbf{v}) k(\mathbf{v}) d\mathbf{v} \tag{3.4}$$

for $\mathbf{v} = \mathbf{s}_{i'} - \mathbf{u}$ and $\mathbf{d} = \mathbf{s}_i - \mathbf{s}_{i'}$.

Higdon et al. (1998) propose letting the kernel function vary with spatial location:

$$V(\mathbf{s}) = \int k(\mathbf{s} - \mathbf{u}; \mathbf{s}) z(\mathbf{u}) d\mathbf{u} \tag{3.5}$$

For model fitting they parameterize the kernel function and estimate the parameters using a hierarchical Bayesian approach. A scale parameter is used to control the size of the spatially varying kernels and hence the posterior distribution of this parameter is used to investigate how the range of the spatial dependence changes over the spatial domain. Higdon et al. (1998) relied heavily on specifying $k(.)$ based on a Gaussian kernel. Paciorek and Schervish (2004) offer a larger class of kernels including a spatially varying Matérn kernel.

Higdon (1998) propose a space–time model based on a convolution of a finite sum approximation to (3.5) so that space–time process $V(\mathbf{s}, t)$ is represented as a weighted average of a smooth white noise process $z(\mathbf{s}, t)$ across a set of discreet space–time coordinates, $\mathbf{X}_l = (\mathbf{x}_l, \tau_l)$, $l = 1, \ldots, M$. Let $z(\mathbf{X}_l)$ be iid N(0,1) for $l = 1, \ldots, M$.

$$V(\mathbf{s}, t) = \sigma_j \sum_{l=1}^{M} k(\mathbf{s} - \mathbf{x}_l, t - \tau_l; \mathbf{s}) z(\mathbf{x}_l, \tau_l) \tag{3.6}$$

Gelfand et al. (2004) suggest modeling p processes, $V_j(\mathbf{s})$, $j = 1, \ldots, p$ by using different kernel functions $k_j(\mathbf{s} - \mathbf{u}; \mathbf{s})$ for each process following (3.5). This is immediately extended to the space–time scenario based on the model in (3.6).

$$V_j(\mathbf{s}, t) = \sum_{l=1}^{M} k_j(\mathbf{s} - \mathbf{x}_l, t - \tau_l; \mathbf{s}) z(\mathbf{x}_l, \tau_l) \tag{3.7}$$

Now the cross–covariance for $\mathbf{V}$ is

$$Cov(V_j(\mathbf{s}_i, t_k), V_{j'}(\mathbf{s}_{i'}, t_{k'})) = \sigma_j \sigma_{j'} \sum_{l=1}^{M} k_j(\mathbf{s}_i - \mathbf{x}_l, t_k - \tau_l; \mathbf{s}_i) k_{j'}(\mathbf{s}_{i'} - \mathbf{x}_l, t_{k'} - \tau_l; \mathbf{s}_{i'}) \tag{3.8}$$

Notice that a change in the order of the variables changes the value of the cross–covariance function, i.e. $Cov(V_j(\mathbf{s}_i, t_k), V_{j'}(\mathbf{s}_{i'}, t_{k'})) \neq Cov(V_{j'}(\mathbf{s}_i, t_k), V_j(\mathbf{s}_{i'}, t_{k'}))$. This feature of the cross–covariance function is due to the fact that $k_j(.; \mathbf{s}) \neq k_{j'}(.; \mathbf{s})$ and is not necessarily related to the cross–covariance of the actual multivariate process.

A drawback of this approach is that each component of the multivariate process is a mixture of the same white noise process, imposing strong correlation between the variables. Without loss of generality consider the case when M=1 and $\mathbf{s}_i = \mathbf{s}_{i'}$ for fixed time t_k.

$$Corr(V_j(\mathbf{s}_i), V_{j'}(\mathbf{s}_i)) = \frac{\sigma_j \sigma_{j'} K_j(\mathbf{s}_i - \mathbf{x}_1; \mathbf{s}_i) K_{j'}(\mathbf{s}_i - \mathbf{x}_1; \mathbf{s}_i)}{\sqrt{\sigma_j^2 K_j^2(\mathbf{s}_i - \mathbf{x}_1; \mathbf{s}_i)} \sqrt{\sigma_{j'}^2 K_{j'}^2(\mathbf{s}_i - \mathbf{x}_1; \mathbf{s}_i)}} = 1 \tag{3.9}$$

3.2.2 Multivariate Mixture Model

Gelfand et al. (2004) describe an extension of the nonstationary univariate modeling approach of Fuentes and Smith (2001). Using this approach each process is modeled as a convolution of an uncountable number of locally stationary processes.

$$V_j(\mathbf{s}) = \int k_j(\mathbf{s}-\mathbf{x}) z_{\theta(\mathbf{x})}(\mathbf{s}) dx \tag{3.10}$$

$z_{\theta(\mathbf{x})}$, $x \in \mathcal{D}$ is a family of independent stationary Gaussian processes indexed by parameter $\theta(\mathbf{x})$ which varies across space. Note that (3.10) is a stochastic integral and is defined as a limit in mean square of approximating sums, e.g Cressie (1993), p. 107. Let $\Sigma(\mathbf{s}_i - \mathbf{s}_{i'}; \theta(\mathbf{x}))$ be the stationary covariance function for the mean zero Gaussian processes $z_\theta(\mathbf{s})$. For example Fuentes and Raftery (2005) use a Matérn covariance model. In this case $\theta(\mathbf{x}) = (\nu(\mathbf{x}), \sigma^2(\mathbf{x}), \rho(\mathbf{x}))$ for smoothness parameters ν, sill parameters σ^2 and range parameters ρ. If $\theta(\mathbf{x})$ varies slowly across different locations $\mathbf{x}$ and the kernel fuction k_j is sharply peaked then for two sites $\mathbf{s}_i$ and $\mathbf{s}_{i'}$ that are close together $V_j(\mathbf{s}_i)$ and $V_j(\mathbf{s}_{i'})$ behaves like an observation from a stationary process with parameters $\theta(\mathbf{s})$. Fuentes and Smith (2001) refer to this as "local stationarity".

Thus we have

$$Cov(V_j(\mathbf{s}_i), V_{j'}(\mathbf{s}_{i'})) = \int K_j(\mathbf{s}_i - \mathbf{x}) K_{j'}(\mathbf{s}_{i'} - \mathbf{x}) \Sigma(\mathbf{s}_i - \mathbf{s}_{i'}; \theta(\mathbf{x})) dx \tag{3.11}$$

However the limitation with this model is that the spatial covariance function, Σ, is the same for every pair of spatial processes, $V_j(\mathbf{s})$, $V_{j'}(\mathbf{s})$. In practice we approximate the convolution with a discrete set of support points $\{\mathbf{x}_l\}$, $l = 1, \ldots, M$ so that:

$$Cov(V_j(\mathbf{s}_i), V_{j'}(\mathbf{s}_{i'})) = \frac{1}{M} \sum_{l=1}^{M} K_j(\mathbf{s}_i - \mathbf{x}_l) K_{j'}(\mathbf{s}_{i'} - \mathbf{x}_l) \Sigma(\mathbf{s}_i - \mathbf{s}_{i'}; \theta(\mathbf{x}_l)) \tag{3.12}$$

Again if we consider the case when M=1 and $\mathbf{s}_i = \mathbf{s}_{i'}$, it can immediately be shown that $Corr(V_j(\mathbf{s}_i), V_{j'}(\mathbf{s}_i)) = 1$. Hence we see this model, similar to the kernel mixture approach of the previous section, imposes strong correlation between the variables. In section 3.3 we adopt a similar mixture model for the space–time covariance and use the linear model of coregionalization to model multiple variables with differing degrees of correlation.

3.2.3 Spatially Varying LMC

We now turn to the extension of the LMC model offered by Gelfand et al. (2004) for non-stationary data. They propose a Spatially Varying LMC (SVLMC) model based on allowing the weight matrix $\mathbf{A}$ to be a function of spatial location.

$$\mathbf{V}(\mathbf{s}) = \mathbf{A}(\mathbf{s})\mathbf{w}(\mathbf{s}) \tag{3.13}$$

At each location $\mathbf{s}$ the covariance matrix for the p components is now a function of space. For $\mathbf{A}(\mathbf{s})$ a deterministic function we have:

$$Cov(\mathbf{V}(\mathbf{s}), \mathbf{V}(\mathbf{s})) = \mathbf{A}(\mathbf{s})\mathbf{A}(\mathbf{s})^T = \mathbf{T}(\mathbf{s}) \tag{3.14}$$

Gelfand et al. (2004) propose letting $\mathbf{T}(\mathbf{s}) = g(x(\mathbf{s}))\mathbf{T} = (x(\mathbf{s}))^{\psi}\mathbf{T}$ for $x(\mathbf{s})$ some covariate for $\mathbf{V}$ and $\psi \geq 0$ unknown. It is standard practice to model the mean function of a spatial process with a deterministic function (such as a polynomial function) of spatial covariates (such as location or elevation). This version of the SVLMC uses such covariates to model a spatially varying sill parameter. In their commercial real estate example, Gelfand et al. (2004) use a set of covariates for modeling both the mean and the variance in this fashion.

Utilizing the one–to–one relationship between $\mathbf{T}$ and $\mathbf{A}$, we have $\mathbf{A}(\mathbf{s}) = \sqrt{g(x(\mathbf{s}))}\mathbf{A}$. The cross–covariance function has the form:

$$Cov(\mathbf{V}, \mathbf{V}) = \sum_{j=1}^{p} \mathbf{T}_j \otimes \mathbf{R}_j^* \tag{3.15}$$

where $\mathbf{R}_j^*$ are a set of p $n \times n$ matrices with $(R_j^*)_{ii'} = \sqrt{g(x(\mathbf{s}_i))}\sqrt{g(x(\mathbf{s}_{i'}))}\rho_j(||\mathbf{s}_i - \mathbf{s}_{i'}||; \theta_j)$ and ρ_j are a set of stationary correlation functions.

The second proposal made by Gelfand et al. (2004) is to model $\mathbf{T}(\mathbf{s})$ as a "matrix–variate spatial process" by defining properties for a Spatial Wishart distribution. (See Appendix B.) Now the covariance matrix is itself a random spatial process, $\mathbf{T}^{-1}(\mathbf{s}) \sim SW_p(\nu, \mathbf{\Gamma}\mathbf{\Gamma}^T, \phi_1, \ldots, \phi_p)$.

We investigate what this implies for the distribution of $\mathbf{V}(\mathbf{s})$. To examine the form of the expected value of $\mathbf{V}(\mathbf{s})$ we need the distribution for $\mathbf{A}(\mathbf{s})$. Let $\mathbf{Z}_{p\times\nu}$ be a matrix with elements $\mathbf{Z}_{jl}(\mathbf{s})$ indep N(0,$\phi_j(\mathbf{s}-\mathbf{s}')$) for $l = 1\ldots\nu$, $j = 1,\ldots p$. Consider this as p different spatial processes with ν replications of each. (Note, ϕ_j denotes the spatial correlation functions for each process.) The Spatial Wishart distribution is defined as:

$$\mathbf{\Omega}_{p\times p}(\mathbf{s}) = \mathbf{\Gamma}\mathbf{Z}(\mathbf{s})\mathbf{Z}^T(\mathbf{s})\mathbf{\Gamma}^T \sim SW_p(\nu, \mathbf{\Gamma}\mathbf{\Gamma}^T, \phi_1, \ldots, \phi_p) \tag{3.16}$$

Letting $\mathbf{T}^{-1}(\mathbf{s}) = \mathbf{\Omega}(\mathbf{s})$ and solving for the elements of the weight matrix:

$$\mathbf{A}(\mathbf{s}) = \begin{pmatrix} (\Omega_{11}(\mathbf{s}) - \Omega_{12}^2(\mathbf{s}))^{-\frac{1}{2}} & 0 \\ -\Omega_{12}(\mathbf{s})(\Omega_{11}(\mathbf{s})\Omega_{22}(\mathbf{s}) - \Omega_{12}^2(\mathbf{s})\Omega_{22}(\mathbf{s}))^{-\frac{1}{2}} & \Omega_{22}^{-\frac{1}{2}}(\mathbf{s}) \end{pmatrix} \tag{3.17}$$

As an alternative Gelfand et al. (2004) mention modeling the weight matrix directly with Gaussian processes. The example given by Gelfand et al. (2004) is to let $a_{jj'}(\mathbf{s}) = z_{jj'}(\mathbf{s})$; $j \neq j'$ and $a_{jj}(\mathbf{s}) = |z_{jj}(\mathbf{s})|$; for $z_{jj'}(\mathbf{s})$ independent mean 0 Gaussian processes for $1 \leq j' \leq j \leq p$. For example for $p = 2$,

$$\mathbf{A}(\mathbf{s}) = \begin{pmatrix} |z_{11}(\mathbf{s})| & 0 \\ z_{21}(\mathbf{s}) & |z_{22}(\mathbf{s})| \end{pmatrix} \tag{3.18}$$

Now the covariance matrix takes the form:

$$\mathbf{T}(\mathbf{s}) = \begin{pmatrix} z_{11}^2(\mathbf{s}) & |z_{11}(\mathbf{s})|z_{21}(\mathbf{s}) \\ |z_{11}(\mathbf{s})|z_{21}(\mathbf{s}) & z_{21}^2(\mathbf{s}) + z_{22}^2(\mathbf{s})| \end{pmatrix} \tag{3.19}$$

Their argument against this route is that this is now a "nonstandard and computationally intractable distribution" for $\mathbf{T}(\mathbf{s})$. However by the same token we argue that the expression in (3.17) for the distribution of the weight matrix $\mathbf{A}(\mathbf{s})$ is also nonstandard and difficult to interpret. Gelfand et al. (2004) suggest this as an advantage of the method stating that the multivariate process "is not only nonstationary but nonGaussian". It is unclear how to decide when this nonGaussian distribution would be appropriate for a given application.

For prediction purposes the SVLMC offers a flexible framework but for applications in geostatistics we are also interested in being able to interpret the estimated parameters. We thus propose to avoid these issues altogether by modeling the latent processes $\mathbf{w}$ in the LMC model as nonstationary space–time processes. This yields a nonstationary multivariate space–time model for $\mathbf{V}(\mathbf{s})$ based on a paramerization that is easier to derive and interpret. For example in the following section we shall see how the nonstationary LMC model allows us to analyze how the parameters of the spatial correlation change in different regions of the hurricane surface wind fields. There is also a single parameter used to measure the strength of the correlation between wind components.

3.3 A New Nonstationary Multivariate Model

3.3.1 Spatial–Temporal LMC

We propose an intuitive multivariate space–time framework by modeling the univariate latent processes as spatial–temporal processes. Thus we specify a set of spatial–temporal covariance models for $\mathbf{w}(\mathbf{s}, t)$ that are allowed to be nonseparable and nonstationary. Extending the linear model of coregionalization approach outlined in the previous chapter:

$$\tilde{\mathbf{V}}(\mathbf{s}, t) = \mathbf{A}\mathbf{w}(\mathbf{s}, t) \tag{3.20}$$

where w_j are a set of $j = 1, \ldots, p$ independent processes each with mean zero and $n \times n$ space–time covaraince matrix:

$$Cov(\mathbf{w}_j, \mathbf{w}_j) = \mathbf{R}_j(\theta_j) \tag{3.21}$$

for θ_j the parameter(s) of the covariance function. At this stage we define $\mathbf{R}_j$ as a set of nonstationary kernel mixture functions proposed by Fuentes et al. (2005). The space–

time structure is defined through a series independent space–time processes in a set of M subregions of stationarity $\mathbf{S}_1, \ldots, \mathbf{S}_M$. Let $\mathbf{X}_l$ represent the centroid of the lth subregion, $\mathbf{X}_l = (\mathbf{s}_l, t_l)$, $l = 1, \ldots, M$. Within each subregion we define the stationary space–time covariance function $C_j(\mathbf{s}_i - \mathbf{s}_{i'}, t_k - t_{k'}; \theta_j(\mathbf{X}_l))$. The space–time dependency may be separable or nonseparable. For example, Fuentes et al. (2005) use a nonseparable parametric model based on a 3D Matérn type covariance. The covariance function for each w_j is the kernel mixture:

$$(R_j)_{uu'} = \rho_j(\mathbf{s}_i - \mathbf{s}_{i'}, t_k - t_{k'}) = \sum_{l=1}^{M} k(\mathbf{s}_i - \mathbf{s}_l) k(\mathbf{s}_{i'} - \mathbf{s}_l) C_j(\mathbf{s}_i - \mathbf{s}_{i'}, t_k - t_{k'}; \theta_j(\mathbf{X}_l)) \quad (3.22)$$

The cross–covariance matrix for $\mathbf{V}(\mathbf{s}, t)$ can be written as for the original LMC model:

$$Cov(\tilde{\mathbf{V}}, \tilde{\mathbf{V}}) = \sum_{j=1}^{p} \mathbf{T}_j \otimes \mathbf{R}_j \quad (3.23)$$

where $\mathbf{T}_j = a_j a_j^T$ and a_j are the j^{th} columns of matrix $\mathbf{A}$. However this approach changes the original role of the weight matrix $\mathbf{A}$. In the traditional LMC model the underlying w_j processes are assumed to have variance one and ρ_j are defined as correlation functions accordingly. In this case there is a one–to–one relationship between the $\mathbf{A}$ matrix and the cross–covariance matrix $\mathbf{T}$. In other words, the sill parameters for each V_j process is a function of these weights, e.g. $Var(V_1) = a_{11}^2$, $Var(V_2) = a_{21}^2 + a_{22}^2$, etc. We use a nonstationary *covariance* matrix that allows the sill parameter to be a function of spatial location. In this case the $\mathbf{A}$ matrix is only used to define the cross–covariance between the different variables. Recall we define $\mathbf{A}$ to be lower triangular. Now we add the condition that $A_{jj} = 1$ for $j = 1, \ldots, p$ leaving $(p^2 - p)/2$ elements to be estimated. For example for $p = 2$:

$$\mathbf{A} = \begin{pmatrix} 1 & 0 \\ a_{21} & 1 \end{pmatrix} \quad (3.24)$$

and thus the nonstationary nonseparable full cross–covariance can be written in the form:

$$Cov(\mathbf{V},\mathbf{V}) = \begin{pmatrix} \mathbf{R}_1 & a_{21}\mathbf{R}_1 \\ a_{21}\mathbf{R}_1 & a_{21}^2\mathbf{R}_1 + \mathbf{R}_2 \end{pmatrix} \tag{3.25}$$

Note that when the range parameters of the covariance functions, ρ_j, are all equal, this does not create an intrinsic model, such as was introduced in Chapter 2 with the standard LMC approach. The variance of the random fields, $w_j(\mathbf{s}, t)$, (the sill parameters) change over the spatial domain and as a result the correlation between two variables will still be a function of space in this case.

3.3.2 Statistical Methodology

A Bayesian framework is used for estimation and prediction of the hurricane wind fields based on the nonstationary space–time model. The Blocking Gibbs Sampling algorithm outlined in Chapter 2 is used to estimate the parameters of the Holland mean function, bias and measurement error in the observed data as well as the parameters of the spatial–temporal covariance. The priors for the random effects, bias and mean function parameters follow what was outlined in Chapter 2. Noninformative priors are used for the spatial range and sill parameters for each of the M subregions of stationarity. The range parameters are assigned uniform priors based on the size of the subregions. The sill parameters are assigned inverse gamma priors. The weight parameter a_{21} is assigned a $N(0, .001)$ prior (.001 is the precision). The next section describes the spatial–temporal covariance model. A separable covariance is found to be appropriate for this application. The temporal covariance is modeled as an autoregressive model and the priors for the coefficients of the AR model are specified as $N(0, .01)$ (.001 is the precision).

For the nonstationary mixture model we use a two dimensional kernel function, $K(\mathbf{s} -$

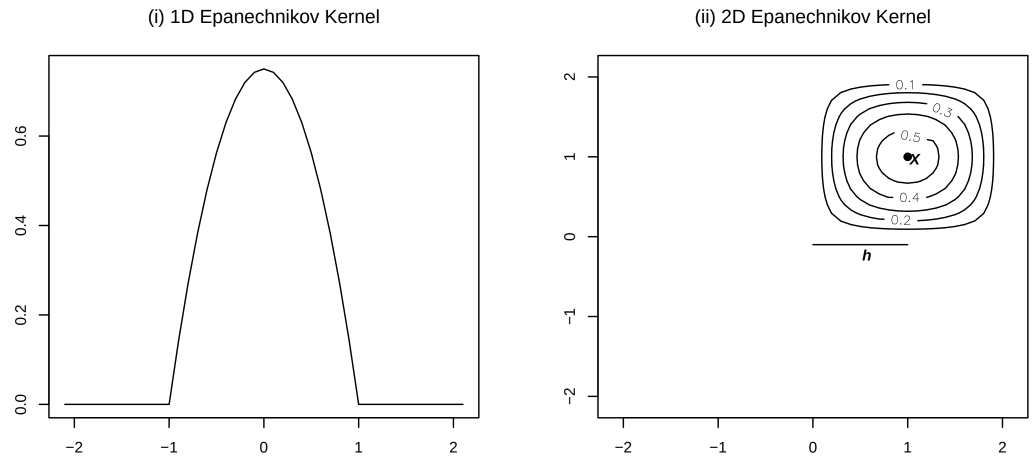

Figure 3.1: (i) Example of 1D Epanechnikov kernel, (ii) example of 2D Epanechnikov kernel centered at support point $\mathbf{x}$ with bandwidth h.

$\mathbf{x}_l) = \frac{1}{h_l^2} K_0(\frac{\mathbf{s}-\mathbf{x}_l}{h_l})$ where K_0 is the two dimensional Epanechnikov kernel (see Figure 3.1(ii)).

$$K_0(\mathbf{u}) = \frac{3}{4}(1-u_1^2)_+ \frac{3}{4}(1-u_2^2)_+ \tag{3.26}$$

The Schwartz Bayesian Information Criterion (BIC) is used to select the number of subregions of stationarity, M (Fuentes et al., 2005). We divide the gridded HRD data into two regions and estimate the covariance parameters of nonstationary LMC model. We continue to divide one subregion at a time in half (e.g. 3 subregions, 4 subregions etc.) until there is no significant improvement in the parameter estimation based on the BIC values. The support points, $\mathbf{x}_l$, are defined as the center of the lth subregion and the bandwidth, h_l is defined as half the maximum distance for the lth subregion.

3.4 Analysis of Data from Hurricane Floyd

We apply the multivariate spatial–temporal model for hurricane surface winds to a case study of Hurricane Floyd. Three different statistical models are compared through a series of model diagnostics. The posterior predictive distribution is used to predict the winds at a set of buoy locations and cross validation is used to evaluate the prediction results.

Hurricane Floyd moved up the Eastern United State's coast September 15th – 16th, 1999. It made landfall on September 16th near Cape Fear, North Carolina and moved over the Eastern part of the state causing massive inland flooding. Floyd was a large and intense hurricane with measured winds close to the threshold for a category 5 storm. The wind speed contours (m/s) for Floyd based on HRD gridded data are shown in Figure 3.2. Unlike Hurricane Charley, the focus of our analysis in Chapters 1 and 2, Floyd had a much larger radius of maximum winds. As a result, the residual wind fields (observed winds minus the Holland wind field) show large residual velocities over a larger spatial domain.

Looking at HRD data from Sept. 15th 730 UTC to 2230 UTC, the residual winds are largest to the right of the hurricane's center. Figure 3.3 shows the residual u winds in the four quadrants of the storm for September 15th, 730 UTC. The Holland mean function consistently under–estimates the velocities in the two right–hand quadrants. The sample variance is also higher in these regions suggesting different levels of spatial variability in different regions of the spatial domain. Exploratory analysis based on sample variograms for data in different quadrants of the storm also provided evidence of spatial nonstationarity in these residual wind fields.

The HRD wind fields from six time periods, 730 UTC, 1030 UTC, 1330 UTC, 1630 UTC, 1930 UTC and 2230 UTC, are used to estimate the nonstationary LMC model. Ten minute average wind speed and direction data is available from the National Data Buoy Center. The buoy data is used to estimate bias in the HRD fields. This data set also provides data during time periods when no new HRD data is available.

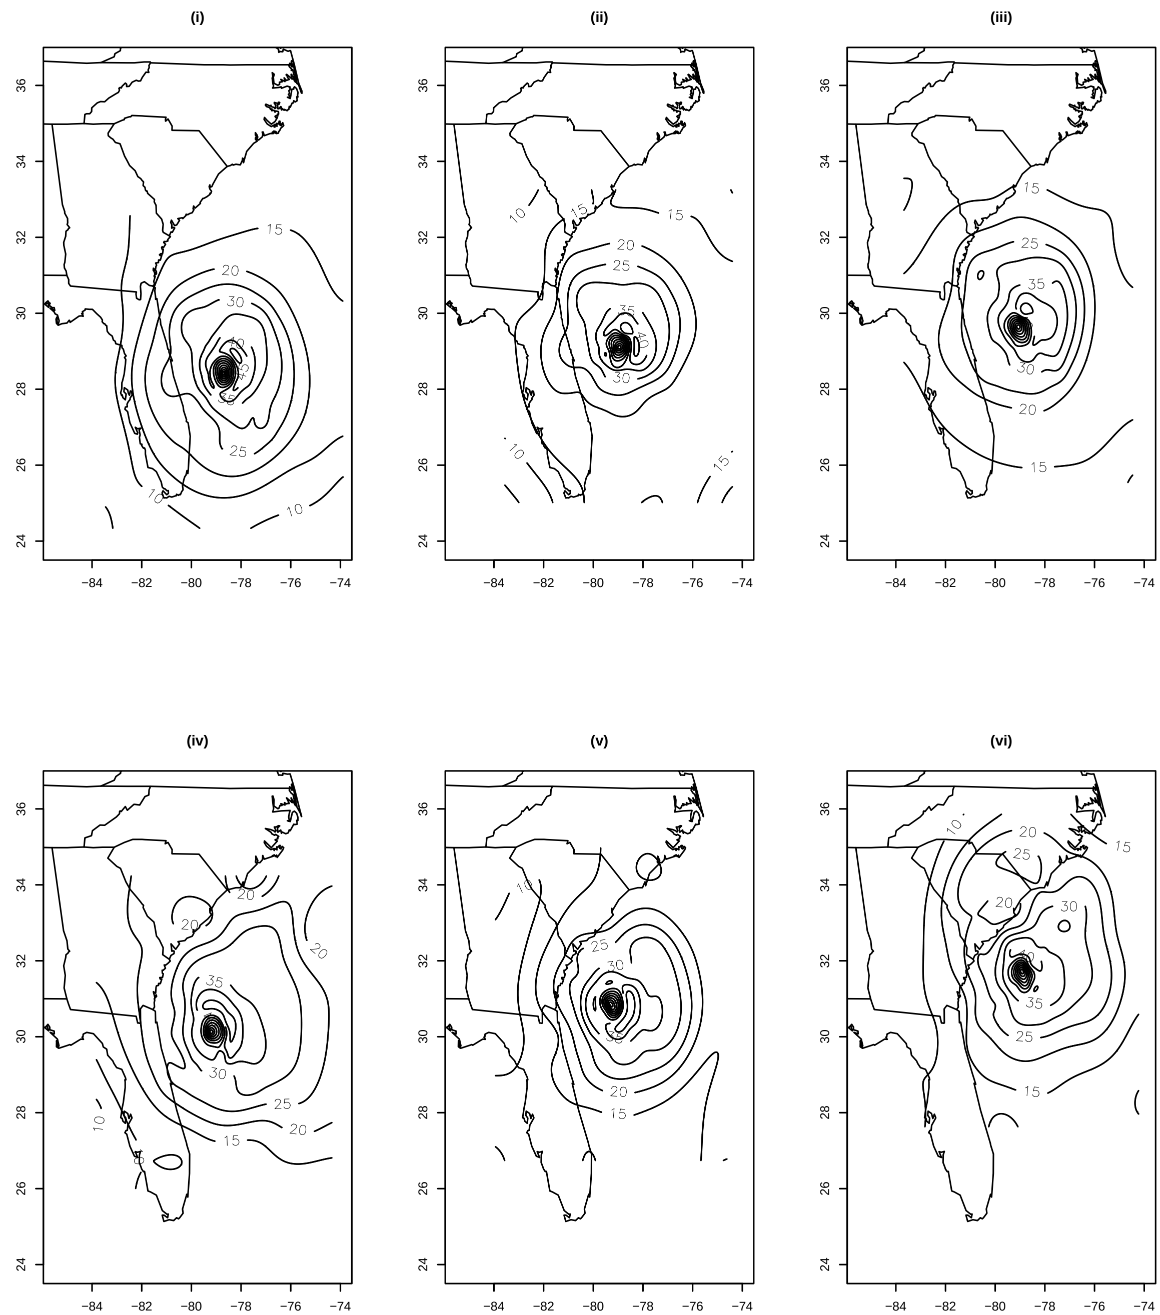

Figure 3.2: Wind contours (m/s) for September 15th hour (i) 730 UTC, (ii) 1030 UTC, (iii) 1330 UTC, (iv) 1630 UTC, (v) 1930 UTC, (vi) 2230 UTC.

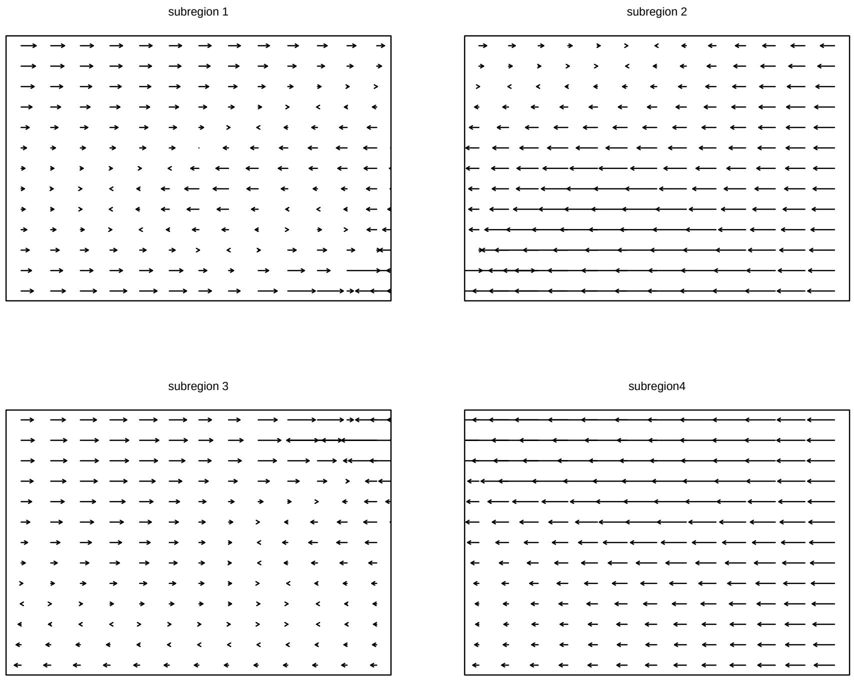

Figure 3.3: Residual u winds (observed – estimated) in four subregions.

3.4.1 Model Estimation for HRD and Buoy Data

In this section we describe the space–time model used to improve estimation of the wind fields for Hurricane Floyd. The HRD wind observations were found to be temporally independent once the Holland mean is subtracted. Recall that the HRD winds are at three hour intervals. Between these times the strength and structure of the hurricane can change drastically. We propose to treat these three hour periods as windows of spatial–temporal stationarity. A space–time model is fit using all available buoy and HRD data within a three hour window. Furthermore the space–time structure is allowed to change between different windows. Note that this was also the approach adopted in Chapter 2 to account for changes in the spatial structure of Hurricane Charley.

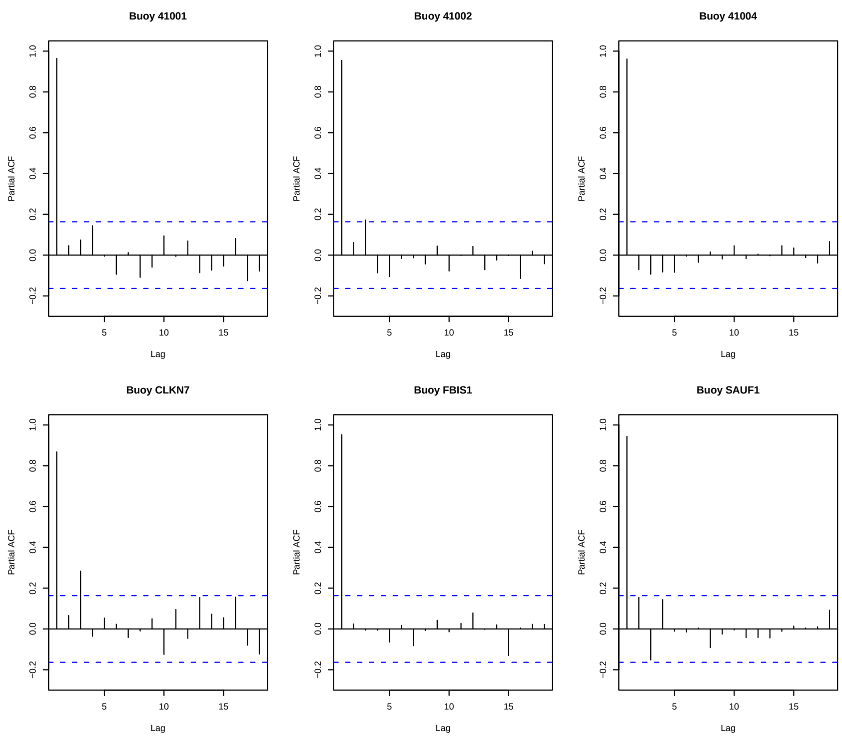

Figure 3.4: Partial autocorrelation for v winds at 6 buoy locations. The time interval is 10 minutes.

The time series plots of the residual u and v winds are very similar at the set of buoy locations. Figure 3.4 shows the PACF plots for the residual v winds at six sites. We also look at the spatial correlation between the buoys at different times. Overall there is not strong evidence of a spatial–temporal interaction in the buoy residual wind data. The covariance functions, C_j, are thus modeled as separable spatial–temporal functions with a Gaussian spatial covariance matrix and a temporal covariance corresponding to an autoregressive model of order one.

The Gaussian model (also called the squared exponential), $\rho(d;\phi_j) = exp(-(d/\phi)^2)$, for $d = ||\mathbf{s}_i - \mathbf{s}_{i'}||$, is infinitely differentiable which induces a very smooth spatial process. We note that in many applications it is unrealistic to assume the unknown covariance function of the observed data is so highly differentiable. Using the AIC and BIC criterion we compare the exponential, Mat´ern and Gaussian covariance models and find that the Gaussian model was the best fit for the re–analysis surface wind fields provided by the HRD. In fact, this results makes sense based on how these winds are derived. The HRD wind fields are a combination of surface weather observations from ships, buoys, coastal platforms, surface aviation reports, reconnaissance aircraft data, and geostationary satellites which have been processed to create gridded data a a very high spatial resolution (approximately 6km $\times$ 6km with the number of grid points on the order of 10^4 for each time period). Thus although the original buoy or ship data may be more spatially heterogeneous, it is realistic to represent the output fields of the HRD re–analysis winds as a very smooth spatial process.

The BIC criterion is used to choose the number of regions of stationarity (M). We find that dividing the spatial domain into four subregions corresponding to the four quadrants surrounding the storm center is sufficient. Thus the subregions of spatial stationarity change for each three hour window based on the location of the hurricane center provided by the HRD data.

Figure 3.5 shows the posterior distributions for the sill and range parameters for the u winds for 730 UTC. The densities are in fact significantly different in the different subregions. The sill parameters are larger in subregions 2 and 4 which corresponds to the empirical analysis based on the sample variances. The residual winds are more homogeneous in region 4, as seen in Figure 3.4, and the estimated range parameter is also largest in this region. Thus we see that the nonstationary LMC approach allows for empirical validation of the model since the parameters can still be interpreted in the geostatistical sense.

Model Diagnostics and Calibration

Three different models are fit based on the HRD and buoy data at each of the six time periods. Model 1 is the stationary LMC model used in Chapter 2. Model 2 uses $M = 4$ subregions of stationarity and treats the u and v components as independent. Model 3 is

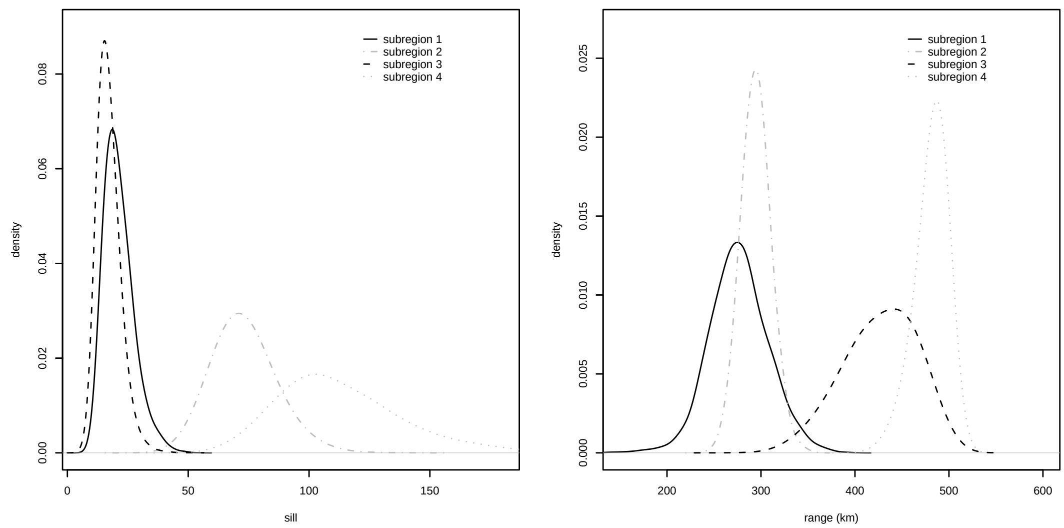

Figure 3.5: Posterior distributions for (i) the sill parameter, (ii) the range parameter for 700 UTC.

the full nonstationary space time LMC proposed here. The Bayesian Information Criterion (BIC) (Schwarz, 1978) for each model is shown in Table 3.1. The posterior median is used for the parameter estimates. The nonstationary models, Models 2 and 3, show a consistent improvement over Model 1. As we saw in Chapter 2, the difference between the independent model and multivariate model is less pronounced. Still, except for hours 1630 and 1930 UTC, using the LMC model provides a better fit to the data.

Table 3.1: BIC for each time period using Models 1, 2 and 3.

Model	730 UTC	1030 UTC	1330 UTC	1630 UTC	1930 UTC	2230 UTC
Model 1	8117.2	14399.0	9030.7	7861.9	17543.0	6079.4
Model 2	5867.3	5998.8	5706.5	5492.0	5885.6	5709.1
Model 3	5775.8	5901.9	5474.9	5757.3	6059.7	5639.3

For cross validation of the three models, the median values of the posterior predictive distributions are used to compare to the observed values at thirty HRD grid points that were not used for parameter estimation. The root means square prediction error (RMSPE) values are shown in Table 3.2. The values are very similar but the nonstationary models do show an improvement over the stationary model at all times except for hour 730 UTC.

Table 3.2: RMSPE for $\mathbf{V}$ (m/s) based on 30 cross validation sites using Models 1, 2 and 3.

Model	730 UTC	1030 UTC	1330 UTC	1630 UTC	1930 UTC	2230 UTC
Model 1	1.56	3.60	3.91	1.56	4.41	3.38
Model 2	1.97	2.28	3.31	1.29	2.73	2.87
Model 3	1.67	2.35	2.90	1.28	2.81	2.65

3.4.2 Wind Field Prediction

The Rao–Blackwellized estimator for the full posterior predictive distribution $P(\mathbf{V}(\mathbf{s}_o, t_o)|\hat{\mathbf{V}})$ is used to predict the wind fields at location $\mathbf{s}_o$ and time t_o. The full posterior predictive distribution is estimated for every ten minute interval between September 15th, 730 UTC and 2230 UTC. Figure 3.6 shows the observed and predicted v velocities for two buoys off the coast of South Carolina. (Similar results were found for the u component.) The dashed and dotted line is the output from the Holland model using the posterior median values for B and $Rmax$. Even using the parameter estimates rather than fixed values, the Holland function cannot capture the change in the winds as the hurricane approaches these locations. The posterior predicted values (dashed lines) for Model 3 (plots i-a and ii-a) are almost identical to the observed data at each location. The spatial and temporal correlation is able to capture the variability in the wind velocity. The corresponding 95% posterior predictive intervals are also included as dotted lines. The predicted values from the stationary model, Model 1 (plots i-b and ii-b), are also very close to the observed values but the observed wind values fall outside of the 95% intervals 13 out 90 times for Buoy 41004.

For cross validation we predict the v winds at the buoy locations without using the

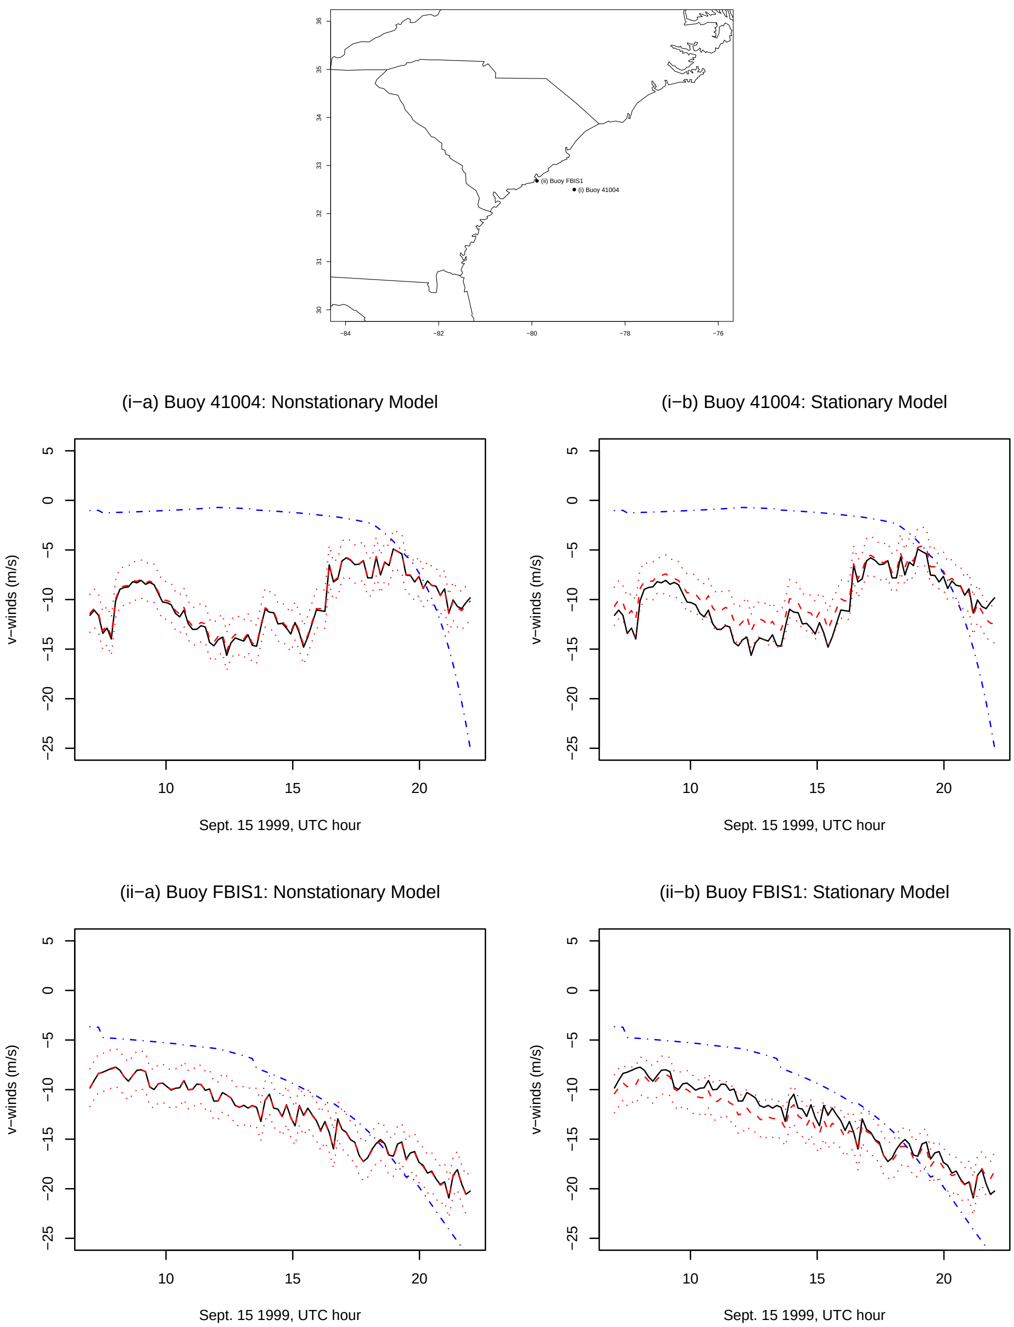

Figure 3.6: Map of 2 NBDC buoys. The remaining plots show the observed v winds in m/s at each buoy in solid black for September 15th 730 UTC to 2230 UTC. The dashed and dotted line is the Holland output based on posterior median values for B and $Rmax$. The dashed lines are the predicted values and the dotted lines are 95% posterior predictive intervals using (i-a, ii-a) Model 3 and (i-b, ii-b) Model 1.

observations at these sites. Figure 3.7 shows the predicted values under Model 3 (plots i-a and ii-a) and Model 1 (plots i-b and ii-b). The nonstationary model is still able to better capture the variability in the winds when compared to the stationary LMC model. The predicted values are closer to the Holland mean since the posterior predictive distribution is conditioned on less spatial–temporal information. The 95% posterior prediction intervals for the nonstationary model are wider, accounting for more uncertainty in the estimates. Although the nonstationary model predicts the v winds reasonably well for buoy FBIS1, the calibration percent for buoy 41004 is only 63%. However the calibration of the stationary model estimates is much lower; for both buoys the observed winds fall inside the 95% intervals only 40% of the time.

3.5 Discussion

The full nonstationary space time modeling techniques for multivariate data proposed here show promising results for hurricane surface wind field prediction. Since the parameterization of the model is built upon standard geostatistical and time series models the parameter estimates can be compared to simple empirical analysis as an initial validation of the results. More rigorous comparisons are needed to quantify the improvement of this method over other approaches outlined by Gelfand et al. (2004).

The methods proposed here rely heavily on the assumption that the data is normally distributed. This is an appropriate assumption for the residual wind fields but may not apply in other applications. The SVLMC offered an approach for nonGaussian data. We note that the kernel mixture model used here can also be extended to the nonGaussian case. This is achieved by letting the local stationary processes, $z_{\theta(\mathbf{x})}$, be a family of nonGaussian processes. Currently we only work with the kernel mixture through the covariance function since the integral in equation (3.10) is defined as a limit in the mean square sense. For the

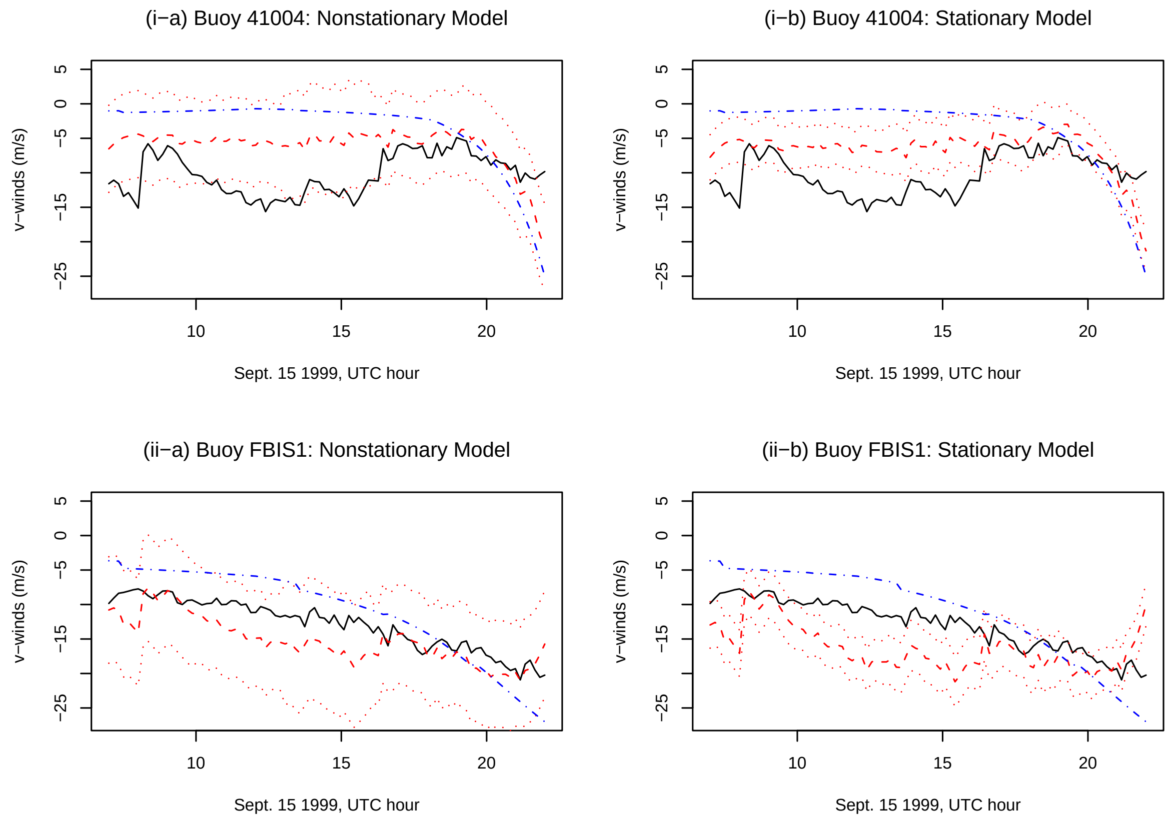

Figure 3.7: Cross Validation Time Series Plots: Observed v winds in m/s at each buoy in solid black for September 15th 730 UTC to 2230 UTC. The dashed and dotted line is the Holland output. The dashed line is the predicted values and the dotted lines are 95% credible intervals under (i-a, ii-a) Model 3 and (i-b, ii-b) Model 1.

Gaussian case only the first two moments are needed but more development of this approach is necessary to define the moments of this distribution for the nonGaussian case.

For the nonstationary kernel mixture model we chose the number of subregions by dividing the HRD spatial domain into rectangles. Other approaches are possible. For example Fuentes et al. (2005) use a k–means clustering approach. Also the bandwidth parameter, h_l, of the kernel functions was fixed for our analysis. Estimation results can be sensitive to the choice of bandwidth so it would also be possible to put a uniform prior distribution on this parameter and estimate it as part of the Gibbs sampling algorithm.

In terms of the scientific motivation of this study, the Bayesian framework offers an attractive way to combine different sources of data and to account for bias and measurement error in these observations. In this way wind data and water level observations were utilized to improve output from the numerical ocean model. This is the first time that these methods have been applied to an ocean model application to be used for improved storm surge forecasting. The numerical ocean model used in this study has been developed specifically for modeling along the Eastern United States coast. Unlike global and other large scale forecasting models, this model can be applied at a higher resolution to a much smaller regional domain. Using statistical methods to combine observed data and model output at this smaller scale is very challenging but it will ultimately have the greatest impact on the health and safety of citizens living along the coast.

Bibliography

Banerjee, S., Carlin, B. P. and Gelfand, A. E. (2004). *Hierarchical modeling and analysis for spatial data*. Chapman & Hall. New York.

Berliner, L. M., Royle, J. A., Wikle, C. K. and Milliff, R. F. (1999). Bayesian methods in the atmospheric sciences. *Bayesian Statistics 6 – Proceedings of the Sixth Valencia International Meeting* pp. 83–100.

Blumberg, A. F. and Mellor, G. L. (1987). A description of a three-dimensional coastal ocean circulation model. *in* N. S. Heaps, ed., 'Three-Dimensional Coastal Ocean Models'. American Geophysical Union. Washington, D.C.. pp. 1–16.

Brown, P. J., Le, N. D. and Zidek, J. V. (1994). Multivariate spatial interplation and exposure to air pollutants. *Canadian Journal of Statistics*.

Canizares, R., Heemink, A. W. and Vested, H. J. (1996). Sequential data assimilation in a fully non-linear hydrodynamic model. *in* A. Muller, ed., 'Hydroinformatics '96'. Balkema. Rotterdam. pp. 463–470.

Canizares, R., Heemink, A. W. and Vested, H. J. (1998). Application of advanced data assimilation methods for the initialization of storm surge models. *Journal of Hydraulic Research* **36**, 655–674.

Cressie, N. (1993). *Statistics for Spatial Data*. Wiley. New York.

Cripps, E., Nott, D., Dunsmuir, T. M. and Wikle, C. K. (2005). Space-time modelling of Sydney Harbour winds. *Australian and New Zealand Journal of Statistics* **47**, 3–17.

De Iaco, S. and D. E. Myers, D. E. D. P. (2003). The linear coregionalization model and the product-sum space-time variogram. *Mathematical Geology* **35**, 25–38.

Echevin, V., De Mey, P. and Evensen, G. (2000). Horizontal and vertical structure of the representer functions for sea surface measurements in a coastal circulation model. *Journal of Physical Oceanography* **30**, 2627–2635.

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research* **99**, 10143–10162.

Evensen, G. (2003). The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics* **53**, 343–367.

Evensen, G. and van Leeuwen, P. J. (1996). Assimilation of geosat altimeter data for the Agulhas Current using the ensemble Kalman filter with a quasigeostrophic model. *Monthly Weather Review* **124**, 85–96.

Ezer, T. and Mellor, G. L. (1994). Continuous assimilation of geostat altimeter data into a three-dimensional primitive equation Gulf Stream model. *Journal of Physical Oceanography* **24**, 832–847.

Ezer, T. and Mellor, G. L. (1997). Data assimilation experiments in the Gulf Stream region: How useful are satellite-derived surface data for nowcasting the subsurface fields?. *Journal of Atmospheric and Oceanic Technology* **14**, 1379–1391.

Fuentes, M. and Raftery, A. (2005). Model evaluation and spatial interpolation by Bayesian combination of observations with outputs from numerical models. *Biometrics* **61**, 36–45.

Fuentes, M. and Smith, R. L. (2001). A new class of nonstationary spatial models. *Tech report, North Carolina State University, Raleigh.*

Fuentes, M., Chen, L., Davis, J. M. and Lackmann, G. M. (2005). Modeling and predicting complex space-time structures and patterns of coastal wind fields. *Environmetrics* **16**, 449–464.

Gelfand, A. E. and Ghosh, S. K. (1998). Model choice: A minimum posterior predictive loss approach. *Biometrika* **85**, 1–11.

Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**, 398–409.

Gelfand, A. E., Schmidt, A. M., Banerjee, S. and Sirmans, C. F. (2004). Nonstationary multivariate process modeling through spatially varying coregionalization (with discussion). *TEST* **13**, 1–50.

Heemink, A. W., Bolding, K. and Verlaan, M. (1997). Storm surge forecasting using Kalman filtering. *Journal of the Meteorological Society of Japan* **75**, 305–318.

Higdon, D. (1998). A process-convolution approach to modeling temperatures in the North Atlantic Ocean. *Journal of Environmental and Ecological Statistics* **5**, 173–190.

Higdon, D., Swall, J. and Kern, J. (1998). Non-stationary spatial modeling. *in* J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith, eds, 'Bayesian Statistics 6'. Oxord University Press. pp. 761–768.

Holland, G. (1980). An analytic model of the wind and pressure profiles in hurricanes. *Monthly Weather Review* **108**, 1212–1218.

Houston, S. H. and Powell, M. D. (1994). Observed and modeled wind and water-level response from tropical storm Marco (1990). *Weather and Forecasting* **9**, 427–439.

Houston, S. H., Shaffer, W. A., Powell, M. D. and Chen, J. (1999). Comparisons of HRD and SLOSH surface wind fields in hurricanes: Implications for storm surge modeling. *Weather and Forecasting* **14**, 671–686.

Hsu, S. A. and Yan, Z. (1998). A note on the radius of maximum wind for hurricanes. *Journal of Coastal Research* **14**, 667–668.

Ide, K., Courtier, P., Ghil, M. and Lorenc, A. C. (1997). Unified notation for data assimilation: Operational, sequential and variational. *Journal of the Meteorological Society of Japan* **75**, 181–189.

Jones, J. E. and Davies, A. M. (1998). Storm surge computations for the Irish Sea using a three-dimensional numerical model including wave-current interaction. *Continential Shelf Research* **18**, 201–251.

Kalnay, E. (2003). *Atmospheric Modeling, Data Assimilation and Predictability.* Cambridge University Press. Cambridge.

Kibria, B. M. G., Sun, L., Zidek, J. V. and Le, N. D. (2002). Bayesian spatial prediction of random space-time fields with application to mapping $pm_{2.5}$ exposure. *Journal of the American Statistical Association* **97**, 112–124.

Large, W. G., Morzel, J. and Crawford, G. B. (1995). Accounting for surface wave distortion of the marine wind profile in low-level Ocean Storms wind measurements. *Journal of Physical Oceanography* **25**, 2959–2971.

Le, N. D. and Zidek, J. V. (1992). Interpolation with uncertain spatial covariance: A Bayesian alternative to kriging. *Journal Multivariate Analysis* **43**, 351–374.

Le, N. D., Sun, W. and Zidek, J. V. (1997). Bayesian multivariate spatial interpolation with data missing by design. *Journal of the Royal Statistical Society, Series B* **59**, 501–510.

Luong, B., Blum, J. and Verron, J. (1998). A variational method for the resolution of a data assimilation problem in oceanography. *Inverse Problems* pp. 979–997.

Madsen, H. and Canizares, R. (1999). Comparison of extended and ensemble Kalman filters for data assimilation in coastal area modelling. *International Journal for Numerical Methods in Fluids* **31**, 961–981.

Mardia, K. V., Kent, J. T. and Bibby, J. M. (1979). *Multivariate Analysis.* Academic Press. London.

Meinhold, R. J. and Singpurwalla, N. D. (1983). Understanding the Kalman filter. *The American Statistician* **37**, 123–127.

Mellor, G. L. and Ezer, T. (1991). A Gulf Stream model and an altimetry assimilation scheme. *Journal of Geophysical Research* **96**, 8779–8795.

Paciorek, C. J. and Schervish, M. J. (2004). Nonstationary covariance functions for Gaussian process regression. *Advances in Neural Information Processing Systems* **16**, 273–280.

Peng, M., Xie, L. and Pietrafesa, J. (2004). A numerical study of storm surge and inundation in the Croatan-Albemarle-Pamlico Estuary System. *Estuarine, Coastal and Shelf Science* **59**, 121–137.

Pietrafesa, L. J., Xie, L., Morrison, J., Janowitz, G. S., Pelissier, J., Keeter, K. and Neuherz, R. A. (1997). Numerical modelling and computer visualization of the storm surge in and around the Croatan-Albemarle-Pamlico Estuary system produced by Hurricane Emily of August 1993. *Mausam* **48**, 567–578.

Powell, M. D. and Houston, S. H. (1996). Hurricane Andrew's landfall in South Florida. Part II: Surface wind fields and potential real-time applications. *Weather Forecast* **11**, 329–349.

Powell, M. D., Houston, S. H. and Reinhold, T. A. (1996). Hurricane Andrew's landfall in South Florida. Part I: Standardizing measurements for documentation of surface wind fields. *Weather Forecast* **11**, 304–328.

Royle, J. A., Berliner, L. M., Wikle, C. K. and Milliff, R. (1999). A hierarchical spatial model for constructing wind fields from scatterometer data in the Labrador Sea. *Case Studies in Bayesian Statistics Volume IV* pp. 367–382.

Schmidt, A. M. and Gelfand, A. E. (2003). A Bayesian coregionalization approach for multivariate pollutant data. *Journal of Geophysical Research* **108**, 8783.

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics* **6**, 461–464.

Spiegelhalter, D. J., Best, N. G., Calin, B. P. and van der Linde, A. (2002). Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B* **64**, 583–640.

Spiegelhalter, D., Thomas, A., Best, N. and Gilks, W. (1996). *Bugs .5 Bayesian inference using Gibbs sampling. Manual, version ii.* MRC Biostatistics Unit, Institute of Public Health. Cambridge, U.K.

Stein, M. L. (1999). *Interpolation of Spatial Data. Some Theory for Kriging.* Springer. New York, N.Y.

Sun, W., Le, D., Zidek, J. V. and Burnett, R. (1998). Assessmetn of a Bayesian multivariate interpolation approach for health impact studies. *Environmetrics* **9**, 565–586.

Verlaan, M. and Heemink, A. W. (1997). Tidal flow forecasting using reduced rank square root filters. *Stochastic Hydrology and Hydraulics* **11**, 349–368.

Wackernagel, H. (1998). *Multivariate geostatistics: An introduction with applications.* Springer. New York.

Wikle, C. K., Berliner, L. M. and Cressie, N. (1998). Hierarchical Bayesian space-time models. *Environmental and Ecological Statistics* **5**, 117–154.

Wikle, C. K., Milliff, R. F., Nychka, D. and Berliner, L. M. (2001). Spatiotemporal hierarchical Bayesian modeling: Tropical ocean surface winds. *Journal of the American Statistical Association* **96**, 382–397.

Xie, L., Bao, S., Pietrafesa, L., Foley, K. and Fuentes, M. (2006). A real-time hurricane surface wind forecasting model: Formulation and verification. *Monthly Weather Review* **134**, 1355–1370.

Xie, L., Pietrafesa, L. and Peng, M. (2004). Incorporation of a mass-conserving inundatoin scheme into a three dimensional storm surge model. *Journal of Coastal Research* **20**, 1209–1223.

Yanglom, A. M. (1962). *An Introduction to the Theory of Stationary Random Functions*. Dover Publications. New York.

Appendices

Appendix A

Multivariate Normal Result

For Gaussian prior and likelihood distributions we can use the following result for conditional distributions (e.g. see Mardia et al. (1979), pp. 63).

- If $\begin{pmatrix} X_1 \\ X_2 \end{pmatrix} \sim N\left[\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \right]$,
  then $(X_1 | X_2 = x_2) \sim N(\mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(x_2 - \mu_2), \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21})$.

- Conversely, if $(X_1 | X_2 = x_2) \sim N(\mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(x_2 - \mu_2), \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21})$ and $X_2 \sim N(\mu_2, \Sigma_{22})$, then $\begin{pmatrix} X_1 \\ X_2 \end{pmatrix} \sim N\left[\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \right]$

Appendix B

Spatially Varying Wishart Distribution

B.1 Wishart Distribution

The Wishart distribution is the multivariate analog of the χ^2 distribution. For matrix $\mathbf{Z}_{p\times\nu}$ with elements $\mathbf{Z}_{jl}$ iid N(0,1), $l = 1\ldots\nu$, $j = 1,\ldots p$ and scale matrix $\mathbf{\Gamma}_{p\times p}$:

$$\mathbf{\Omega}_{p\times p} = \mathbf{\Gamma}\mathbf{Z}\mathbf{Z}^T\mathbf{\Gamma}^T \sim W_p(\nu, \mathbf{\Gamma}\mathbf{\Gamma}^T) \tag{B.1}$$

- $E(\Omega) = \nu\Gamma\Gamma^T$

B.2 Spatially Varying Wishart Distribution

Gelfand et al. (2004) propose a spatially varying version of the Wishart distribution. In this case they define $\mathbf{Z}_{p\times\nu}$ as a matrix with elements $\mathbf{Z}_{jl}(\mathbf{s})$ indep N(0,$\phi_j(\mathbf{s}-\mathbf{s}')$) for $l = 1\ldots\nu$, $j = 1,\ldots p$. Consider this as p different spatial processes with with ν replications of each. (Note, ϕ_j denotes the spatial correlation functions for each process.)

$$\mathbf{\Omega}_{p\times p}(\mathbf{s}) = \mathbf{\Gamma}\mathbf{Z}(\mathbf{s})\mathbf{Z}^T(\mathbf{s})\mathbf{\Gamma}^T \sim SW_p(\nu, \mathbf{\Gamma}\mathbf{\Gamma}^T, \phi_1, \ldots, \phi_p) \tag{B.2}$$

When $\mathbf{\Gamma}$ is diagonal we have the following properites for $\mathbf{\Omega}(\mathbf{s})$ (Gelfand et al., 2004):

$$\begin{aligned}
Cov(\Omega_{jj}(\mathbf{s}), \Omega_{jj}(\mathbf{s}')) &= 2\nu\Gamma^4_{jj}\phi^2_j(\mathbf{s}-\mathbf{s}') \\
Cov(\Omega_{jj'}(\mathbf{s}), \Omega_{jj'}(\mathbf{s}')) &= 2\nu\Gamma^2_{jj}\Gamma^2_{jj'}\phi_j(\mathbf{s}-\mathbf{s}')\phi'_j(\mathbf{s}-\mathbf{s}') \qquad j \neq j' \\
Cov(\Omega_{jj'}(\mathbf{s}), \Omega_{kk'}(\mathbf{s}')) &= 0 \qquad\qquad j \neq k \text{ or } j' \neq k'
\end{aligned} \tag{B.3}$$

For example, for $p = 2$ and $\mathbf{\Gamma} = \text{diagonal}(\gamma_1, \gamma_2)$:

$$\mathbf{\Omega}(\mathbf{s}) = \mathbf{\Gamma}\mathbf{Z}(\mathbf{s})\mathbf{Z}^T(\mathbf{s})\mathbf{\Gamma}^T = \begin{pmatrix} \gamma_1^2 \sum_{l=1}^{\nu} Z_{1l}^2(\mathbf{s}) & \gamma_1\gamma_2 \sum_{l=1}^{\nu} Z_{1l}(\mathbf{s})Z_{2l}(\mathbf{s}) \\ \gamma_1\gamma_2 \sum_{l=1}^{\nu} Z_{1l}(\mathbf{s})Z_{2l}(\mathbf{s}) & \gamma_2^2 \sum_{l=1}^{\nu} Z_{2l}^2(\mathbf{s}) \end{pmatrix} \tag{B.4}$$

- $E(\mathbf{\Omega}(\mathbf{s})) = \begin{pmatrix} \nu\gamma_1^2 & 0 \\ 0 & \nu\gamma_2^2 \end{pmatrix}$
- $Cov(\Omega_{11}(\mathbf{s}), \Omega_{11}(\mathbf{s}')) = 2\nu\gamma_1^4\phi_1^2(\mathbf{s} - \mathbf{s}')$
- $Cov(\Omega_{12}(\mathbf{s}), \Omega_{12}(\mathbf{s}')) = 2\nu\gamma_1^2\gamma_2^2\phi_1(\mathbf{s} - \mathbf{s}')\phi_2(\mathbf{s} - \mathbf{s}')$
- $Cov(\Omega_{22}(\mathbf{s}), \Omega_{22}(\mathbf{s}')) = 2\nu\gamma_2^4\phi_2^2(\mathbf{s} - \mathbf{s}')$
- $Cov(\Omega_{jj'}(\mathbf{s}), \Omega_{kk'}(\mathbf{s}')) = 0$ for $j \neq k$ or $j' \neq k'$.

Appendix C

WinBUGS Code for Gibbs Sampler

```
model {
# Likelihood
# HRD Data
for(i in 1:nr){
    Ur[i] ~ dnorm(mu.ur[i],gamma.uA)                    # V̂^a = a + bV + ε^a
    mu.ur[i] ← au+bu*Ut[i]
    Ut[i] ]← mu.uH[i] + Sru[i]                          # V = V^H + Ṽ
    Vr[i] ~ dnorm(mu.vr[i],gamma.vA)
    mu.vr ← av+bv*Vt[i]
    Vt[i] ← mu.v[i] + Srv[i]
    mu.v[i] ← mu.vH[i]+alpha*(Ut[i]-mu.uH[i])          # Conditional mean for v|u.
}
# Buoy Data
for(i in (nr+1):n){
    Ur[i] ~ dnorm(Ut[i],gamma.uB)                       # V̂^b = V + ε^b
    Ut[i] ]← mu.uH[i] + Sru[i]                          # V = V^H + Ṽ
    Sru[i] ← mu.u[i]+Ut[i]
    Vr[i] ~ dnorm(Vt[i],gamma.vB)
    Vt[i] ← mu.v[i] + Srv[i]
    mu.v[i] ← mu.vH[i]+alpha*(Ut[i]-mu.uH[i])          # Conditional mean for v|u.
}

# Conditional mean specifications:
for(i in 1:n){
    wH[i] ← sqrt(pow(Rmax,B)*B*(1010*100-Pc*100)*exp(-pow(Rmax,B)/ (pow(radius[i],B)))/
          (1.2*pow(radius[i],B)))                       # Holland function for wind speed.
    mu.uH[i] ← usign[i]*wH[i]*sin(phi[i])              # Orthogonal wind components based
    mu.vH[i] ← vsign[i]*wH[i]*cos(phi[i])              # on axis – symmetric wind field.
    mu[i] ← 0
}

# Latent spatial processes:
Sru[1:n] ~ spatial.exp(mu[ ],x[ ],y[ ],tau.u,p1,kappa)
Srv[1:n] ~ spatial.exp(mu[ ],x[ ],y[ ],tau.v,p1,kappa)
```

```
# Priors:
# Parameters for measurement error
    gamma.uR ← 1/pow(t.uR,2)
    gamma.vR ← 1/pow(t.vR,2)
    t.uR ∼ dunif(.01,3)
    t.vR ∼ dunif(.01,3)
    gamma.uB ← 1/pow(t.uB,2)
    gamma.vB ← 1/pow(t.vB,2)
    t.uB ∼ dunif(.01,1)
    t.vB ∼ dunif(.01,1)
# Parameters for LMC
    p1 ∼ dunif(.001,.1)
    p2 ∼ dunif( .001,.1)
    theta1 ← 1/p1
    theta2 ← 1/p2
    alpha ∼ dnorm(0,.01)
    tau.u ∼ dgamma(.01,.01)
    sigma.u ← 1/sqrt(tau.u)
    tau.v ∼ dgamma(.01,.01)
    sigma.v ← 1/sqrt(tau.v)
# Parameters for (Holland) mean function
    B ∼ dunif(1,2.5)
    Rmax ∼ dunif(5,60)
# Bias parameters
    au ∼ dnorm(0,.01)
    av ∼ dnorm(0,.01)
    bu ∼ dunif(.9,1.1)
    bv ∼ dunif(.9,1.1)
# Additional parameters of interest:
    T[1,1] ← 1/tau.u
    T[1,2] ← alpha/tau.u
    T[2,2] ← pow(alpha*sigma.u,2)+1/tau.v
    T[2,1] ← T[1,2]
    corr ← T[1,2]/(sqrt(T[1,1])*sqrt(T[2,2]))
}
```