

**IMPROVING STORM SURGE HAZARD CHARACTERIZATION USING
“PSEUDO-SURGE” TO AUGMENT HYDRODYNAMIC SIMULATION
OUTPUTS**

by

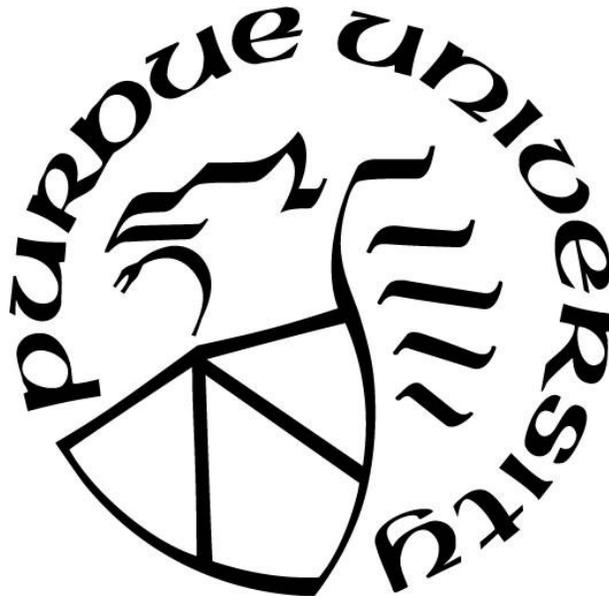
Matthew Shisler

A Thesis

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Master of Science in Industrial Engineering



School of Industrial Engineering

West Lafayette, Indiana

May 2019

**THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL**

Dr. David R. Johnson, Chair

School of Industrial Engineering

Dr. Susan R. Hunter

School of Industrial Engineering

Dr. Venkatesh M. Merwade

Lyles School of Civil Engineering

Approved by:

Dr. Steven J. Landry

Head of the Graduate Program

To friends and family

ACKNOWLEDGMENTS

The author would like to thank:

- The United States Air Force, for selecting me to obtain this degree and the opportunity of continued service to the nation.
- The Purdue Military Research Initiative, for providing the opportunity for me to pursue a degree at Purdue University.
- Dr. David Johnson, for his guidance, counseling, and patience over countless discussions.
- Dr. Susan Hunter, for her technical guidance and assistance
- Dr. Venkatesh Merwade, for his technical guidance and assistance

PREFACE

DISCLAIMER CLAUSE: The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

TABLE OF CONTENTS

LIST OF TABLES	viii
LIST OF FIGURES	ix
LIST OF ABBREVIATIONS.....	xi
ABSTRACT.....	xii
1. INTRODUCTION	1
2. LITERATURE REVIEW	4
2.1 Quantitative Definition of Risk.....	4
2.2 Estimating the Probability Distribution Function for Storm Surge	4
2.2.1 Formulation of Design Storm Events	5
2.2.2 Estimates Based Only on Historical Storms	5
2.2.3 The Joint Probability Method	5
2.2.4 The Empirical Simulation Technique	7
2.2.5 The Empirical Track Method.....	8
2.2.6 Modern JPM with Optimal Sampling.....	9
2.3 Censored Data Techniques	10
3. BACKGROUND AND CONTEXT	13
3.1 Coastal Louisiana Risk Assessment Model	13
3.2 Synthetic Storm Definitions.....	14
3.3 CLARA Geospatial Domain.....	15
3.4 Advanced Circulation and Simulating Waves Nearshore Models.....	17
3.5 CLARA Response Surface Methodology.....	18
4. METHODS	22
4.1 Problem Description	23
4.2 Approaches to Identifying Pseudo-surge Values	27
4.2.1 Topographic Elevation Replacement Rule	27
4.2.2 Combined Regression and Extrapolation	27
4.2.3 Pointwise Selection of Pseudo-surge Values.....	31
4.2.4 Simultaneous Selection of Pseudo-surge Values.....	35
4.2.5 Ad Hoc Approach to Simultaneous Selection of Pseudo-Surge Values.....	39

4.3	Implementation	42
5.	RESULTS & DISCUSSION	44
5.1	Impacts on Response Surface Performance	46
5.2	Impacts on Flood Exceedance Curve Estimates	51
5.3	Discussion	60
6.	CONCLUSIONS AND FUTURE WORK	63
	APPENDIX A. JPM-OS SYNTHETIC STORM DEFINITIONS	65
	APPENDIX B. CLARA WATERSHED CHARACTERISTICS	68
	APPENDIX C. RESPONSE SURFACE METHDOLOGY DISCUSSION	70
	APPENDIX D. ADDITIONAL FLOOD EXCEEDANCE ANALYSIS	72
	REFERENCES	75

LIST OF TABLES

Table 3.1 Summary of Covariates Used in Response Surface Models.....	18
Table 4.1 Storm-Location Set Definitions	25
Table 4.2 Summary Characteristics for Select Watersheds	39
Table 4.3 Ad Hoc Simultaneous Selection of Pseudo-surge Tabulation	41
Table 5.1 Storm Wet/Non-wet Classification Rates (full region).....	48
Table 5.2 Storm Wet/Non-wet Classification Rates (reduced region)	48
Table D.1 Storm Wet/Non-wet Classification Rates (removed heuristic).....	72

LIST OF FIGURES

Fig. 3.1 CLARA Model Module Schematic (Fischbach et al., 2017)	13
Fig. 3.2 Storm Set Definitions	15
Fig. 3.3 Full CLARA Study Region	15
Fig. 3.4 CLARA Grid Points (Fischbach et al., 2017).....	16
Fig. 3.5 CLARA Watersheds (Fischbach et al., 2017)	16
Fig. 3.6 ADCIRC+SWAN Modeling Domain (Roberts & Cobell, 2017).....	17
Fig. 3.7 Observation Density for Simulation Storm Set	19
Fig. 3.8 Observation Density for Grid Points of Interest.....	20
Fig. 3.9 Regression Method by Grid Point	21
Fig. 4.1 Flood Depth Exceedance Probability Estimation Production Workflow	22
Fig. 4.2 Workflow to Generate Baseline True Surge and Flood Depth Exceedances	23
Fig. 4.3 Simple Extrapolation Example.....	28
Fig. 4.4 Example Regression with Pseudo-value Included	29
Fig. 4.5 SSE for Defined Observations as a Function of the Pseudo-value.....	30
Fig. 4.6 Convexity of $f(s)$	34
Fig. 4.7 2-D Locally Weighted Regression Example (Feuillet et al., 2015).....	35
Fig. 4.8 Unknown Impacts of Pointwise Pseudo-surge on LWR	36
Fig. 5.1 Ad Hoc Simultaneous Selection of Pseudo-surge Data Availability	45
Fig. 5.2 Grid Point Wetting RMSE by Observation Density (full region)	46
Fig. 5.3 Grid Point Non-wetting RMSE by Observation Density (full region).....	47
Fig. 5.4 Grid Point Wetting RMSE by Observation Density (reduced region)	49
Fig. 5.5 Grid Point Non-wetting RMSE by Observation Density (reduced region).....	49
Fig. 5.6 Grid Point Wetting RMSE by Original Regression Method	50
Fig. 5.7 Grid Point Non-wetting RMSE by Original Regression Method.....	51
Fig. 5.8 Average Flood Exceedance Estimate Bias by Return Period (<40%) (full region)	52
Fig. 5.9 Flood Exceedance Estimate Bias by Observation Density (100-yr) (full region).....	53
Fig. 5.10 100-year Flood Exceedance Bias Impacts (full region)	54
Fig. 5.11 500-year Flood Exceedance Bias Impacts (full region)	55
Fig. 5.12 Flood Exceedance Estimate Bias by Return Period (<50%) (reduced region).....	56

Fig. 5.13 Flood Exceedance Estimate Bias by Observation Density (100-yr) (reduced region)..	57
Fig. 5.14 100-year Flood Exceedance Bias Impacts (reduced region)	58
Fig. 5.15 500-year Flood Exceedance Bias Impacts (reduced region)	59
Fig. D.1 Grid Point Non-wetting RMSE by Observation Density (removed heuristic)	72
Fig. D.2 Average Flood Exceedance Estimate Bias by Return Period (removed heuristic)	73
Fig. D.3 Flood Exceedance Estimate Bias by Observation Density (100-yr) (rem heuristic).....	74

LIST OF ABBREVIATIONS

AEPF	Annual Exceedance Probability Function
ADCIRC	ADvanced CIRCulation
AHS	Ad Hoc Simultaneous
CDF	Cumulative Distribution Function
CLARA	Coastal Louisiana Risk Assessment
CMP	Comprehensive Master Plan
CPARLWR	Conditionally Parametric Locally Weighted Regression
EAD	Expected Annual Damages
EST	Empirical Simulation Technique
ETM	Empirical Track Method
HURDAT	National HURricane DATabase
JPM-OS	Joint Probability Method with Optimal Sampling
LA CPRA	Louisiana Coastal Protection and Restoration Authority
NAVD88	North American Vertical Datum 1988
NOAA	National Oceanic and Atmospheric Administration
PW	Pointwise
RH	Removed Heuristic
RMSE	Root Mean Squared Error
SPH	Standard Project Hurricane
SQP	Sequential Quadratic Programming
SWAN	Simulated WAVes Nearshore
USACE	U.S. Army Corps of Engineers

ABSTRACT

Author: Shisler, Matthew, P. MSIE

Institution: Purdue University

Degree Received: May 2019

Title: Improving Storm Surge Hazard Characterization Using “Pseudo-Surge” to Augment Hydrodynamic Simulation Outputs

Major Professor: David R. Johnson

Joint probability methods for assessing storm surge flood risk involve the use of a collection of hydrodynamic storm simulations to fit a response surface model describing the functional relationship between storm surge and storm parameters like central pressure deficit and the radius of maximum wind speed. However, in areas with a sufficiently low probability of flooding, few storms in the simulated storm suite may produce surge, with most storms leaving the location dry with zero flooding. Analysts could treat these zero-depth, “non-wetting” storms as either truncated or censored data. If non-wetting storms are excluded from the training set used to fit the storm surge response surface, the resulting suite of wetting storms may have too few observations to produce a good fit; in the worst case, the model may no longer be identifiable. If non-wetting storms are censored using a constant value, this could skew the response surface fit. The problem is that non-wetting storms are indistinguishable, but some storms may have been closer to wetting than others for a given location. To address these issues, this thesis proposes the concept of a negative surge, or “pseudo-surge”, value with the intent to describe how close a storm came to causing surge at a location. Optimal pseudo-surge values are determined by their ability to improve the predictive performance of the response surface via minimization of a modified least squares error function. We compare flood depth exceedance estimates generated with and without pseudo-surge to determine the value of perfect information. Though not uniformly reducing flood depth exceedance estimate bias, pseudo-surge values do make improvements for some regions where <40% of simulated storms produced wetting. Furthermore, pseudo-surge values show potential to replace a post-processing heuristic implemented in the state-of-the-art response surface methodology that corrects flood depth exceedance estimates for locations where very few storms cause wetting.

1. INTRODUCTION

The state of Louisiana and its surrounding coastline repeatedly have been victims of intense tropical storms causing severe damage. In an effort to characterize and counteract flood risk, the Louisiana Coastal Protection and Restoration Authority (CPRA) developed the *Comprehensive Master Plan for a Sustainable Coast* (CMP) (CPRA, 2017). In its most recent iteration, the CMP details 124 restoration, structural and non-structural risk reduction projects to be implemented over the next 50 years to best mitigate the effects of storm hazards and reduce land loss (CPRA, 2017). To make informed decisions, CPRA commissioned the development of a risk assessment model to characterize the level of risk for current and possible future climate conditions. Thus, the Coastal Louisiana Risk Assessment (CLARA) Model was produced for the 2012 release of the LA Coastal Master Plan and later improved for the 2017 CMP revision.

Hazard characterization, a key component of the risk calculation, is in this case an estimate of the annual exceedance probability function (AEPF) for storm surge-based flood depths, commonly referred to as the flood depth exceedance curve. The AEPF is essentially a cumulative distribution function (CDF), but the EPF returns the probability that a given flood depth is met or exceeded in a year timespan. Estimates of the flood depth exceedance curve require a mathematical formulation of storm characteristics and simulation of their effects on the coastal region of interest. The AEPF estimates are important because they impact risk assessment and eventually the CMP project selection process.

Joint probability methods for assessing storm surge flood risk involve the use of a collection of hydrodynamic storm simulations to fit a response surface model describing the functional relationship between storm surge and storm parameters like central pressure deficit and the radius of maximum wind speed. However, in areas with a sufficiently low probability of flooding, few storms in the simulated storm suite may produce surge, with most storms leaving the location dry with no flooding.

Analysts could treat these zero-depth, “non-wetting” storms as either truncated or censored data. If non-wetting storms are excluded from the training set used to fit the storm surge response surface, the resulting suite of wetting storms may have too few observations to produce a good fit; in the worst case, the model may no longer be identifiable. Data truncation with a hierarchy of response surface regression techniques is the current methodology

implemented in the 2017 production version of the CLARA model. On the other hand, censoring non-wetting storms with a constant value could induce substantial bias in wetting storm surge. The problem is that non-wetting storms are indistinguishable, but some storms may have been closer to wetting a location than others (i.e., if the storm was slightly more intense or larger). This is the motivation behind the concept of a negative surge, or “pseudo-surge”, value with the intent to describe how close a storm came to causing surge at a location. Ultimately, the goal of exploring the pseudo-surge concept is not to identify how close an individual storm is to causing surge but to improve predictions of surge from storms outside of a simulated set of storms. The simulated storms and predicted storms are together used to generate flood depth exceedance estimates necessary to characterize the surge hazard.

For a specific location that does not flood during a storm event, one might replace non-wetting observations with pseudo-surge values by implementing a rule that uses information such as 1) the elevation of the non-wet location, 2) distance of the non-wet location to the nearest wet location, and 3) the elevation and positive surge value of the nearest wet location. This thesis completes a key step in the development of such a rule which is to first find optimal pseudo-surge values that best minimize the error of the response surface. The optimal pseudo-surge values are identified with respect to current coastal conditions. A rule of thumb which inserts pseudo-surge values for non-wetting storms would then be applied in future states of the world for which CLARA can generate an assessment of risk.

This thesis proposes and investigates several methods for identifying optimal pseudo-surge values for synthetic storms at non-wetting locations. The performance of the methods is evaluated by comparing the accuracy of surge predictions when the response surface is trained with and without pseudo-surge values. Flood depth exceedance estimates generated with response surface predictions of surge with and without pseudo-surge are compared to gauge the impact on hazard characterization. This demonstrates the potential efficacy of implementing a pseudo-surge value and generate a dataset which can be used to measure the performance of candidate null surge replacement rules.

The optimal pseudo-surge values are found improve flood depth exceedance estimate bias for some regions where fewer than 40% of simulated storms caused wetting. Further, a pseudo-surge value implementation is also shown as a potential replacement for a somewhat

arbitrary response surface post-processing heuristic that is used to reclassify storms as non-wetting.

The thesis is organized in the following manner. Chapter 2 gives a brief review of past literature for estimating the annual exceedance probability function for storm surge and handling censored data in linear regression. Chapter 3 provides background information for the CLARA model and its current role in the 2017 CMP. Chapter 4 describes the methodologies used to select optimal pseudo-surge values that replace undefined surge observations within the construct of the CLARA model framework. Finally, chapters 5 and 6 discuss results, conclusions and future work.

2. LITERATURE REVIEW

2.1 Quantitative Definition of Risk

Kaplan and Garrick (1981) present a quantitative definition of risk as a “set of triplets”: What can happen? How likely is it that that will happen? If it does happen, what are the consequences? In the context of storm surge risk analysis, the first element for which we desire to quantify risk is the tropical storm. The second element is some estimate of the likelihood that a tropical storm will manifest and cause a certain level of storm surge. The final element is a loss function with flood depths as input and damage (usually in monetary terms) as output.

In the case of the flood risk assessment on the Louisiana coastline, a non-trivial approach is used when considering the probability that a location experiences some level of flooding given a storm event (Fischbach et al. 2012). This conditional probability is termed a “vulnerability” and the quantitative definition of risk is defined to be

$$\begin{aligned} \text{risk} &= \text{hazard} \times \text{vulnerability} \times \text{consequence} \\ &= P(\omega) \times P(F_d \geq f_d | \omega) \times L(f_d) \end{aligned}$$

Where ω is a given storm, f_d is a flood depth, and $L(f_d)$ is a loss function that returns damage in dollars due to a flood depth f_d .

This thesis is focused on unprotected coastal areas. The vulnerability term is represented with a point prediction of flood depth due to a specific storm. Characterizing the exceedance probability function for storm surge is equivalent to defining the first two terms of the risk calculation. The flood depth that is associated with losses is generally combined storm surge elevations and wave heights, but here we focus on exclusively on surge. Risk assessment can be enhanced through any of the three terms in the quantitative definition of risk. This thesis focuses on improving the hazard characterization through predictions of flood depths as a function of storm characteristics rather than analyzing the loss function.

2.2 Estimating the Probability Distribution Function for Storm Surge

A number of approaches for estimating the annual exceedance probability function, or a related transformation, have been developed as the practice of flood risk analysis has evolved over time. In the following short sections, we will review some of these methods. The production

version of the CLARA 2017 model (Fischbach et al., 2017) currently implements a modified version of the Joint Probability Method with Optimal Sampling (JPM-OS) (CPRA, 2017).

2.2.1 Formulation of Design Storm Events

Characterizing storm hazards began with the idea of a design storm event. First adopted by the Army Corps of Engineers (U.S. Department of Commerce, 1959; U.S. Weather Bureau, 1965), a “Standard Project Hurricane” (SPH) was meant to be an estimate of a storm that was expected to manifest relatively infrequently for some section of coast. There are two major drawbacks of the design storm event approach. First, historic data was extremely limited at the time the SPH was developed. It so happened that this data was not an accurate representation of typical hurricane activity, at least for the Gulf of Mexico, because of an uncharacteristic lull in storms (Resio, 2009). Second, the design storm was described mathematically with storm intensity as the only degree of freedom. All other storm characteristics were defined as a function of the storm intensity (Resio, 2009). Low variability in the design storm event could lead to highly biased estimates of hurricane impacts.

2.2.2 Estimates Based Only on Historical Storms

The production of an estimate of surge probabilities based solely on historical storms suffers primarily due to the paucity of data (Resio, 2007). Resio argues that if at least one storm is observed each year, then it may be possible to construct such an estimate though one would not be able to apply classic asymptotic statistical methods that assume large sample sizes (Resio, 2007). The fact that we lack detailed records of storms before the mid-20th century and that the frequency of major storms is so low, Resio (2007) further asserts that methods for estimating coastal flooding based solely on historical data should not be used for coastal risk assessment. Another insight not necessarily addressed is the non-stationarity of climate conditions. Using historical storm data, particularly very old data, may produce estimates of flood exceedance probabilities that are no longer relevant to current climate conditions.

2.2.3 The Joint Probability Method

The Joint Probability Method was originally introduced in 1975 in research by the National Oceanic and Atmospheric Administration (NOAA) to define flood risk zones critical in the determination of flood insurance rates and local zoning regulations for the South Carolina

coast (Myers, 1975) and later applied to the Apalachicola and St. George Sound in Florida (Ho & Myers, 1975). Myers describes the JPM to have three main technical aspects: 1) the determination of the climatology of hurricane characteristics, 2) development of a numerical hydrodynamic model to calculate surge levels based on the hurricane atmospheric parameters, and 3) assembling and synthesizing such information into surge frequency analysis (Myers, 1975). These three steps describe the basic approach to storm surge frequency analysis under the JPM.

The first requirement is tackled via the parameterization of a storm with five primary storm characteristics: central pressure, c_p , the radius of the maximum windspeed, r_{max} , the forward velocity of the storm, v_f , the landfall location, x , and the landfall angle, θ . The JPM assumes the structure of each marginal parameter distribution and the joint probability distribution of all five parameters describes the probability of observing a specific storm.

$$f(\omega) = \Lambda_1 \cdot \Lambda_2 \cdot \Lambda_3 \cdot \Lambda_4 \cdot \Lambda_5$$

$$\Lambda_1 = P(c_p|x)$$

$$\Lambda_2 = P(r_{max}|c_p)$$

$$\Lambda_3 = P(v_f|\theta)$$

$$\Lambda_4 = P(\theta|x)$$

$$\Lambda_5 = \psi(x)$$

With $f(\omega) = \Lambda_1 \cdot \Lambda_2 \cdot \Lambda_3 \cdot \Lambda_4 \cdot \Lambda_5$, being the joint probability distribution of storm parameters and $\omega = (c_p, r_{max}, v_f, \theta, x)$ is a five-dimensional vector of storm parameters and $\psi(x)$ is the frequency of storms per year per specified distance along the coast. It is important to note that each marginal parameter distribution is functionally independent of the parameters that are omitted from the conditional statements. For example, since central pressure is conditional with respect to landfall location and the radius of maximum windspeed is conditional with respect to central pressure, then we would expect the radius of maximum windspeed to also be conditional with respect to the landfall location. However, this is not the case because the functional form of the parameter distribution for the radius of maximum windspeed is independent of the landfall location. For more information regarding modern implementations of the JPM, reference Resio (2007) and Resio et al. (2009).

A numerical model, a complex hydrodynamic simulation in modern contexts, is then used to define surge elevation, s , at a location as a function of storm parameters,

$$s(c_p, r_{max}, v_f, \theta, x) = s(\omega).$$

Then the annual rate of occurrence of a surge elevation in excess of η at a site is given by the multiple integral

$$P(s_{max} \geq s) = \lambda \int \dots \int_{\Omega} f(\omega) P(s(\omega) > s) d\omega.$$

Where λ is the historic annual frequency of storms, and $P(s(\omega) > s)$ is the conditional probability of observing a surge elevation greater than s given a storm ω occurs. Since ω fully spans the parameter space, Ω , for synthetic storms, then $P(s_{max} \geq s)$ represents the annual exceedance probability function for storm surge. This is the exceedance curve necessary to define the hazard term in the quantitative definition of risk. In practice, the parameter space Ω is discretized to produce a set of synthetic storms that spans the range of all possible storms. Then an estimate for $P(s_{max} \geq s)$ uses discrete summation in place of the continuous formulation above. The nature of this discretization is further discussed in Sections 2.2.6.

2.2.4 The Empirical Simulation Technique

The Empirical Simulation Technique (EST) is an approach to estimating the flood exceedance probability function pioneered in the early 1990s by Norman Scheffner and Leon Borgman in support of a storm surge frequency analysis conducted by the U.S. Army Corps of Engineers (1999). It was the disadvantages of early versions of the JPM that served as the impetus for the development of EST. The EST was meant to undercut the computational cost of JPM and remove assumptions regarding the structure and mutual independence (though still conditional dependence) of the marginal probability distributions for individual storm parameters. Because the EST removes assumptions regarding the structure of the joint-probability function for the realization of a storm, it is considered a nonparametric approach.

The EST defines an input vector of characteristics, some analogous to those used in the JPM, to describe a storm and samples values for these characteristics with replacement (bootstrapping) from historical storm data to form a training set of storms. Next, a response vector describes the relevant outcomes of a simulation model for a storm event defined by an input vector. Interpolation is then be used to estimate storm outcomes for input vectors not

explicitly modeled. Because the input vector for a given storm is drawn from historical data, “enough” historical data must exist. Furthermore, such sampling implies that a synthetic storm will always match historical data. This is a dangerous assumption when considering future states of the world where environmental and physical infrastructure (e.g., levees, floodwall systems) changes are difficult to predict.

2.2.5 The Empirical Track Method

Both the JPM and the EST require the use of hydrodynamic simulation techniques to record storm surge information. The process of developing a program that can produce data in an accurate and timely manner requires various simplifying assumptions for the modeled environment. Two major assumptions of previous hurricane modeling researchers (Batt et al., 1980; Georgiou et al., 1983; Neumann, 1991; Vickery & Twisdale, 1995) are that a hurricane’s track maintains a straight path or simple curve and that the storm’s central pressure deficit is held constant. The empirical track method developed by Vickery et al. (2000) intends to model the full track of a hurricane or tropical storm. The central pressure is determined as a function of the sea surface temperature and storm heading and speed are updated on a 6-hour schedule (Vickery et al., 2000). The ETM is developed to model hurricane winds but could be extended to model additional storm impacts as well.

The main advantage of this simulation modeling technique is that it allows storm characteristics like direction, intensity, and speed to change during the modeling process. Furthermore, the approach eliminates many problems related to the derivation of hurricane statistics for large model domains. Previous techniques that modeled entire coastlines would fit parameter distributions using all storm data relevant to the domain, smoothing over local climatological features. The empirical track model applies sub-modeling techniques to storm parameters that are based on local data, thus effectively modeling storms for subregions that more closely match the local historic data (Vickery et al., 2000). Such advantages make ETM well-suited for large domains. However, when modeling a subregion of an entire coastline, the Louisiana coast for example, many of the advantages of the ETM are overshadowed by the increased uncertainty associated with fitting probability distributions over time-variant parameters.

2.2.6 Modern JPM with Optimal Sampling

In the mid-2000s two studies compared the efficacy of the JPM and EST performance when estimating flood depth exceedance probability functions (Divoky & Resio, 2007; Agbley & Basco, 2008). Divoky and Resio developed a simulated coast with “hidden rules of nature” to generate a dataset of potentially realistic storm surge observations. They then applied the EST and JPM approaches to this simulated dataset and compared the estimated surge frequencies with the true frequencies implicitly defined by the “hidden rules of nature”. The methods were evaluated by comparing average estimated flood depths to the true flood depths for 50-year, 100-year, and 500-year return periods for several cases with different rules governing simulation replications and the amount of historical data used to build synthetic storm sets. A “return-period”, sometimes referred to as a recurrence interval, describes the estimated annual probability that a flood depth is experienced. For example, a 100-year flood depth is the flood depth which has a 1% chance of being met or exceeded in a given year.

The primary result of the study was that the JPM was found to be robust with respect to assumptions for the structure of the marginal probability distributions for storm parameters even including their mutual dependence. Conversely, EST was found to be highly sensitive to sample variation caused by limited historical data (Divoky & Resio, 2007).

In a similar study, Agbley and Basco (2008) generated a parent population of storms and associated surge responses for a hypothetical coastline and then drew finite sets of storm records to feed the JPM and EST approaches for a comparison. The study produced similar results to Divoky and Resio (2007) determining that EST consistently over-estimated surge and suffered from high variance. Overall, JPM has proven to be more accurate but requires the definition of a large number of synthetic storms. Due to the high computational cost of hydrodynamic simulation models, a high-resolution discretization of the storm parameter space becomes prohibitively expensive. Further research was conducted to optimally sample the storm parameter space and reduce the number of required simulations. These approaches are categorized as Joint Probability Method with Optimal Sampling (JPM-OS).

There are two major methods, a surge response function (Resio, 2007; Resio, 2009; Irish, 2009) and a quadrature method (Toro et al., 2010a) for optimal sampling of the storm parameter space. The quadrature method of Toro et al. selects storms as nodes within a Gaussian quadrature framework in order to maximize the accuracy of a discrete approximation of the exceedance

probability function (2010a). The surge response function is based on the fact that the maximum surge elevation is a reasonably smooth function most sensitive to central pressure, the radius of maximum windspeed and the storm track which allows only a moderate number of simulations from which surge elevations can be interpolated from other storms (Resio, 2007; Resio, 2009; Irish, 2009). Selecting storms in the response surface method is more of a heuristic approach which requires special treatment of storm parameters. These methods were compared against reference data to verify their ability in characterizing surge hazards with a minimum number of synthetic storm simulations (Toro et al., 2010b). Toro found that both approaches yielded appropriate results for application in a flood risk study (2010b). The full details are extensive and not within the scope of this thesis. A modified version of the JPM-OS with response function approach is currently implemented in the CLARA model which recommends the definition of a 446 synthetic storm set.

2.3 Censored Data Techniques

A non-wetting zero-depth storm can be thought of as an example of censored data. Flood depths can be observed as low as zero relative to the topographic elevation but no lower. Modern censored data techniques may prove useful when replacing these non-wetting observations with a meaningful value. The response function used in JPM-OS is assumed to have a linear structure, thus a short overview of some censored data techniques in the context of linear regression is presented in the following section.

Rupert Miller (1976) pioneered efforts in handling censored data in least squares regression in survival analysis of a heart transplant program. In this case, the response variable is survival time of a patient under experimental conditions after the transplant operation. Observations become censored when patients are lost to follow-up during the study for reasons not related to treatment or if some patients are still alive when the study is terminated (Miller, 1976). Therefore, the observations are treated as randomly right-censored (Miller, 1976). In this application Miller used the product-limit estimator of a distribution function, developed by Kaplan and Meier (1958), to assign weights to uncensored observations based on the residuals of a linear model fit with censored and uncensored data and then found regression coefficients that minimize the weighted sum of squared errors for uncensored data (1976).

Extensions to Miller's work have been proposed and refined over time by several biostatisticians. Some notable additions were made by Buckley and James (1979) and Koul et al. (1981). Buckley and James (1979) modified the normal equation which is the analytic solution for multiple linear regression. Koul et al. (1981) developed an estimator with a straightforward extension to multiple linear regression. In a review and comparison of censored data methods, Miller and Halpern (1982) point out that the techniques could be used on left-censored data but maintain the assumption that any censoring is a random event.

Key differences between non-wetting data from a hydrodynamic simulation and censored observations in the aforementioned studies warrant a novel approach. For example, the topographic elevation defines a threshold value below which there is no physical possibility for a lower surge elevation to be observed. Regardless of whether a subject is lost to follow-up or a study terminates in survival analysis, the time until the event of interest will continue to increase. In the case of a storm surge hazard study, when surge estimates fall below a threshold, they simply become null values. Furthermore, the regression methods with censored data assume that the process in which data becomes censored is random whereas surge estimates do not exist below the topographic elevation. Though these studies do not offer methods that can be directly applied to the context of non-wetting in hydrodynamic models, they provide valuable insight when tailoring an approach for selection of optimal pseudo-surge values.

Right-censored data in survival models is the classic case of data censoring, but such issues can also arise in the practice of chemistry. Thompson and Nelson (2003) describe a phenomenon in which chemists falsely report a detectable concentration due to trace amounts of contaminants. Such observations are recorded as "non-detects" or "trace" if some threshold concentration is not satisfied (Thompson & Nelson, 2003). In effort to include these observations in the linear regression estimates, Thompson and Nelson propose using a maximum likelihood approach and compare it to typical substitution rules (2003). This left censoring situation appears more relevant to storm non-wetting than survival analysis but asserts that a "non-detect" implies that the concentration of a containment exists between 0 and the minimum detectable value threshold. A "non-detect" of surge from a hydrodynamic simulation does not share a similar implication. The surge elevation is either measured or it does not exist

This highlights a fundamental difference between the concepts of censored data in the prevailing research and the goal of this thesis. Researchers attempt to impute or consider impacts

from observations which are assumed to have a real value, but pseudo-surge creates a value for observations that lack a real value and serve a singular purpose of improving response function performance.

3. BACKGROUND AND CONTEXT

3.1 Coastal Louisiana Risk Assessment Model

This section will briefly discuss the Coastal Louisiana Risk Assessment (CLARA) model and its role in Louisiana’s Comprehensive Master Plan for a Sustainable Coast. As mentioned earlier, the goal of the CLARA model is to measure risk due to storm surge hazards in terms of expected annual damage (EAD) (in dollars). Additionally, the CLARA model can determine changes in the EAD given the implementation of various wetland restoration or flood risk reduction projects. The CMP consists of a set of projects designed to protect communities from storm hazards and the CLARA model helps evaluate how each project or combination of projects will affect total risk. The flood depth module which models overtopping, system fragility, and interior drainage, has also been used extensively to analyze the flood depth exceedances associated with different projects (Fischbach et al., 2017).

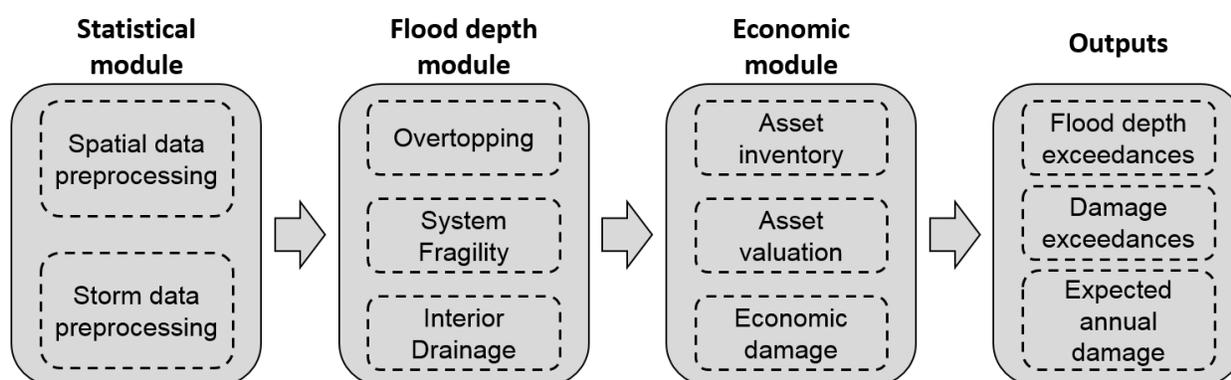


Fig. 3.1 CLARA Model Module Schematic (Fischbach et al., 2017)

Figure 3.1 shows the basic flow of the CLARA model. The JPM-OS procedure, definition of coastal topography and bathymetry, and hydrodynamic simulation of storms occurs in the statistical module which lays the groundwork for the rest of the model operation.

The CLARA model evaluates risk for current coastal conditions (the current state of the world) and many possible states of the future. A scenario is defined with uncertainties about future conditions like storm frequency and intensity, population growth, and fragility of flood mitigation projects. For each scenario CLARA requires flood depth annual exceedance

probability function estimates processed using the results of a hydrodynamic simulation and response surface methodology to calculate EAD. The hydrodynamic simulation of storms is completed separately for each scenario, time period (e.g., 10, 25, or 50 years into the future), and flood mitigation project implementation. State planners then use this information to make informed decisions that will satisfactorily mitigate storm hazards (Fischbach et al., 2017).

3.2 Synthetic Storm Definitions

Historical storm data required in the JPM-OS scheme was sourced from the National Oceanic and Atmospheric Administration's (NOAA) North Atlantic Hurricane database (HURDAT) (Hurricane Research Division, 2014). This data was used to fit marginal distributions for each of the five storm parameters and create the joint distribution for the realization of a storm. Follow-on research was conducted by the US Army Corps of Engineers (USACE) to develop a set of storms which satisfactorily discretized the storm parameter space for use in a coastal Louisiana study (USACE 2008b, USACE 2008c). The result was a set of 446 synthetic storms which we refer to as the "gold standard" set. Because the CLARA model must estimate risk for many possible future states of the world, the computational cost of 446 simulation runs per current and future state is prohibitively expensive. As such, a 92-storm subset of the original 446 storms, referred to as the simulation set of storms, was selected via a bias-efficiency trade-off analysis of the response surface fitting procedure (Fischbach et al., 2017). These 92 storms are used to train the response surface according to the methodology described in Section 3.5. After the response surface function is defined, predictions are made for 76 additional storms (also from the original 446) with track and landfall angles consistent with the storms in the simulation set. These 76 storms will be referred to as the prediction set of storms. For current conditions, simulation data is available for all 446 storms. Simulation and prediction set storm parameters are tabulated in APPENDIX B.

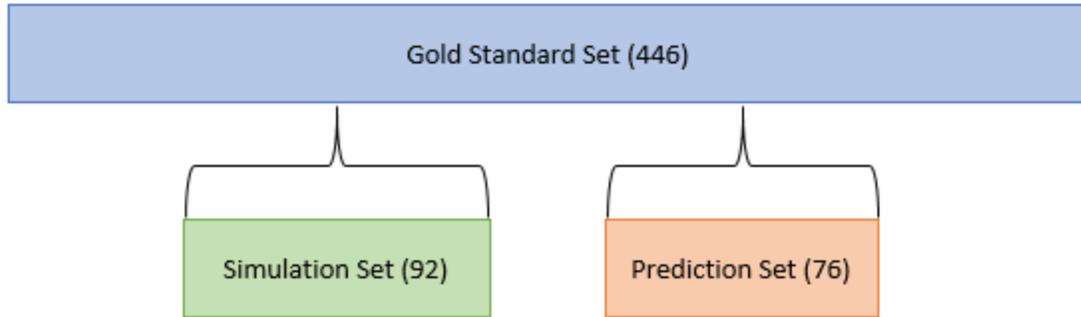


Fig. 3.2 Storm Set Definitions

3.3 CLARA Geospatial Domain

The CLARA geospatial domain defines the study region of the risk assessment model. This domain is then sub-divided in various ways to create meaningful geospatial units of analysis such as state parishes, specialized risk regions, watersheds, census blocks, and grid points (Fischbach et al., 2017). The focus throughout this thesis is on grid points and watersheds.

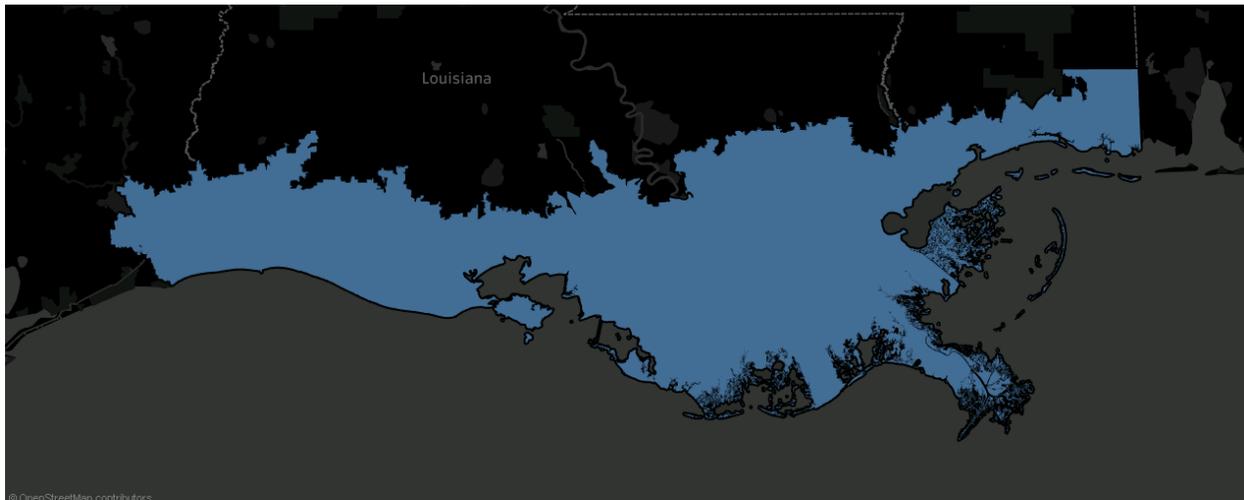


Fig. 3.3 Full CLARA Study Region

The CLARA grid point is the smallest of the geospatial units used in the study (excluding the hydrodynamic simulation mesh). These grid points were selected to balance a tradeoff of spatial fidelity and computational cost. The CLARA grid points are specified through a scripted process which uses the centroids of census blocks and regular spacing when the size of a census block caused a violation of a minimum required 1km resolution. The result is roughly 113,000 grid points that cover the entire study region. Figure 3.4 shows a depiction of these grid points.

Urbanized areas are visible with a high concentration of points and rural areas are marked with regularly spaced points.

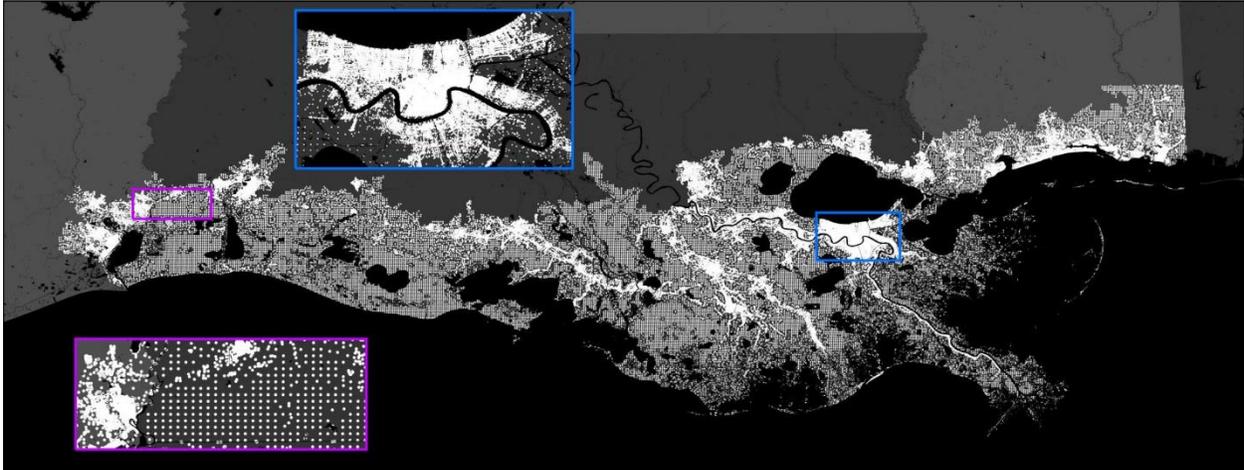


Fig. 3.4 CLARA Grid Points (Fischbach et al., 2017)

A much larger geospatial unit implemented within this study is the watershed. A watershed is a geographic area of independent hydrology. Hydrodynamic activity in one watershed is assumed to have no impact on the activity of another watershed. This may not always be the case, particularly for adjacent watersheds during an extreme tropical storm event. Nonetheless, such definitions are still valuable when generating a response surface from the hydrodynamic simulation results.

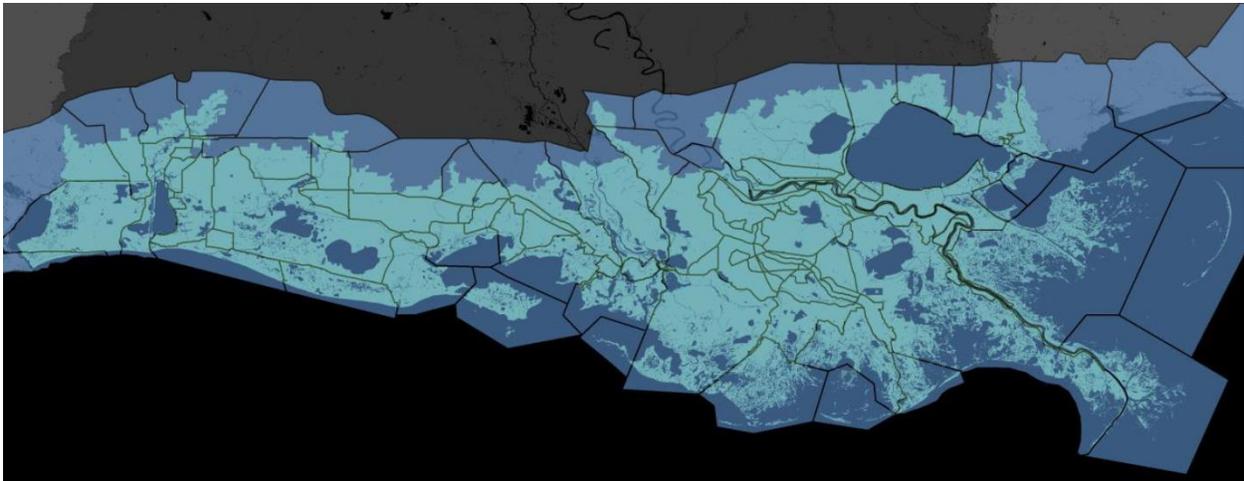


Fig. 3.5 CLARA Watersheds (Fischbach et al., 2017)

The CLARA grid point and watershed are the two relevant geospatial units of analysis when considering augmenting hydrodynamic simulation data in such a way as to improve storm predictions. These units of analysis will be referenced substantially throughout the thesis.

3.4 Advanced Circulation and Simulating Waves Nearshore Models

The ADvanced CIRCulation (ADCIRC) and Simulating Waves Nearshore (SWAN) models fill the role of the hydrodynamic numerical model in the JPM-OS framework implemented in the CLARA statistical module. The ADCIRC+SWAN models generate surge hydrographs as a function of synthetic storm definitions at 1.39 million vertices across the study region and were validated by comparing model outputs to real observations of storms Gustav, Ike, Rita, and Katrina (Roberts & Cobell, 2017).

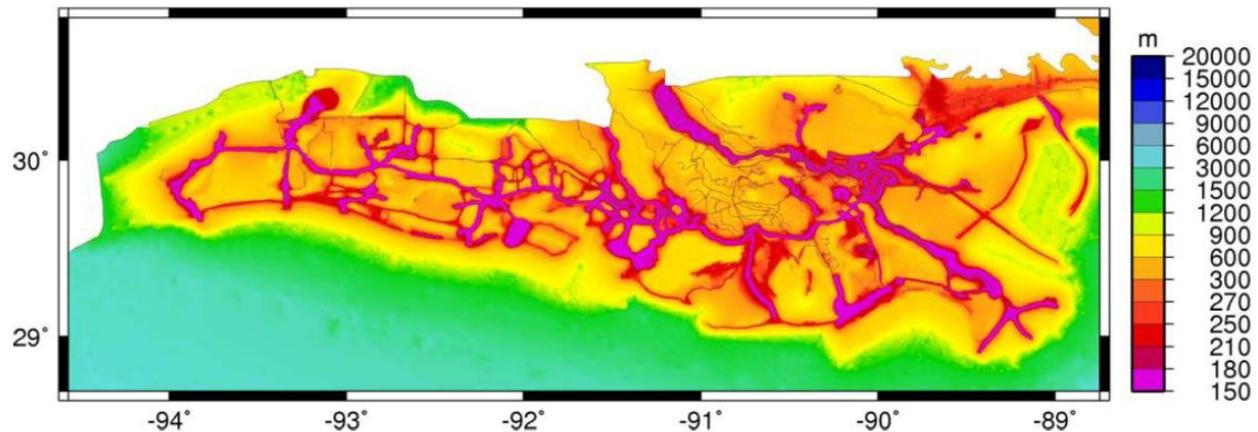


Fig. 3.6 ADCIRC+SWAN Modeling Domain (Roberts & Cobell, 2017)

Figure 3.6 details the model domain where warmer colors indicate areas of higher resolution in the ADCIRC mesh. In areas not enclosed by a storm surge protection system, CLARA receives peak surge elevations at each CLARA grid point which are then used to define the functional relationship between storm surge and storm parameters. Reference Roberts and Cobell (2017), Dietrich et al. (2012) and Luettich & Westerink (2004) for more information regarding the ADCIRC+SWAN models.

3.5 CLARA Response Surface Methodology

This section summarizes details of the state-of-the-art production version of the CLARA 2017 model from Fischbach et al. (2017). The JPM-OS procedure described in Section 2.2.6 outlines the requirements for implementation of a complex hydrodynamic simulation coupled with a response surface function that predicts surge elevations, r_i , relative to the North American Vertical Datum 1988 (NAVD88), at a location i for storms not included in the simulation set. Before the CLARA model calculates an estimate for annual damage due to hurricane hazards, it trains a response surface using the 92-storm simulation set, predicts surge elevations for the 76-storm prediction set, and then uses this information to develop an empirical flood exceedance probability function at each CLARA grid point. In this implementation, non-wetting storms from the hydrodynamic simulation are truncated from the response surface training procedure. The CLARA 2017 production response surface methodology uses the data truncation rule when dealing with non-wetting storms. Both storm characteristics defined in the JPM-OS procedure and geospatial covariates relating the position of a point to the storm track and angle are defined in Table 3.1 below.

Table 3.1 Summary of Covariates Used in Response Surface Models

Covariate	Variable	Units	Description
Central Pressure	c_p	mbars	Minimum atmospheric central pressure
Radius of max windspeed	r_{max}	nm	Lateral size of the storm
Forward velocity	v_f	knots	Speed of the storm
Landfall location	x	-	Landfall location of the storm ¹
Landfall angle	θ	degrees	Angle at which the storm makes landfall ²
Distance	d_{il}	nm	Distance from location i to landfall l
Azimuthal angle	φ_{il}	degrees	Angle between storm track and location i

The response surface is fit using a hierarchy of regression techniques. As data becomes more and more sparse, the procedure reverts to regression techniques which require fewer and fewer covariates. The model hierarchy is as follows:

- (1) Conditionally Parametric Locally Weighted Regression (Cleveland & Devlin, 1988) with full model specification:

¹ Given Louisiana's approximately east-to-west coastline, landfall is defined as the longitudinal point where the storm crosses 29.5 degrees north latitude.

² Relative to the mean angle of historic storms making landfall near a particular landfall location.

$$r_i = \beta_0 + \beta_1 c_p + \beta_2 r_{max} + \beta_3 v_f + \beta_4 d_{il}^3 + \beta_5 d_{il}^2 + \beta_6 d_{il} + \beta_7 \theta_{il} + \beta_8 \sin \varphi_{il} + \beta_9 x + \varepsilon_i$$

(2) Point-by-point ordinary least squares regression with full model specification.

(3) Point-by-point ordinary least squares regression with reduced-form model:

$$r_i = \beta_0 + \beta_1 c_p + \beta_2 r_{max} + \beta_3 v_f + \beta_4 d_{il} + \beta_5 \sin \varphi_{il} + \varepsilon_i$$

(4) Point- and track- level ordinary least squares regression with reduced form model:

$$r_{i,x,\theta} = \beta_0 + \beta_1 c_p + \beta_2 r_{max} + \beta_3 v_f + \varepsilon_i$$

(5) Step function assigning equal surge elevation and wave heights to any synthetic storms more extreme than wetting storms from the simulation set.

The conditions for which the algorithm reduces the complexity of the regression method depends on the number of grid points available and whether a model is identifiable. A model can become unidentifiable if the number of defined observations is less than the number of coefficients to estimate in the model or if there is insufficient variation in the covariate data. Storms in the simulation and prediction sets identified in Fischbach et al. (2017) do not vary in forward velocity. The respective covariate is dropped from the response surface models above.

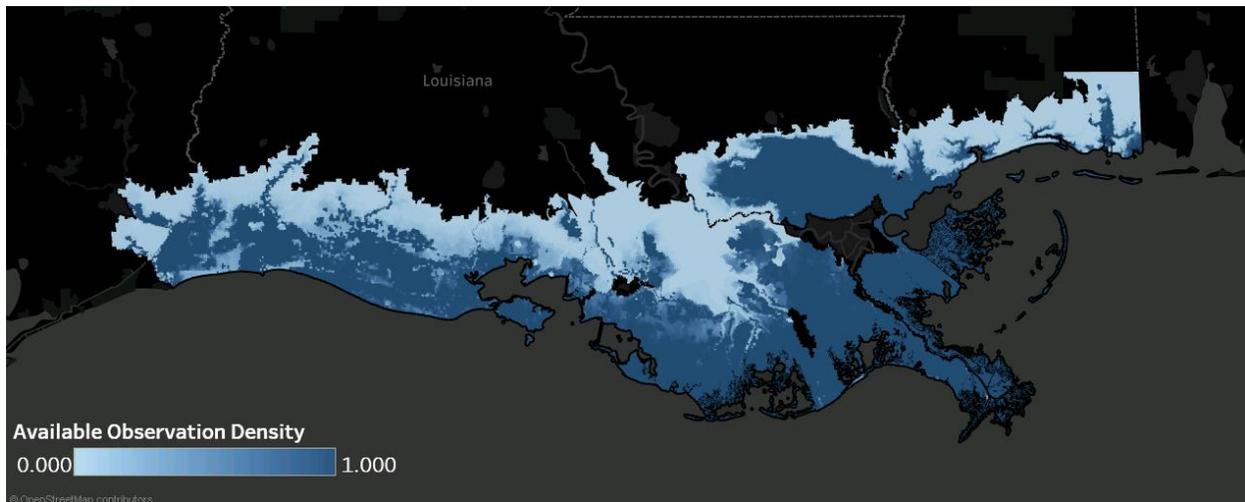


Fig. 3.7 Observation Density for Simulation Storm Set

Figure 3.7 shows the observation density for the simulation storm set under current conditions. The intensity of color represents the proportion of wetting storms in the simulation storm set. Naturally, it is expected that fewer storms produce wetting as distance from the shoreline increases. However, some areas near the shoreline and areas in a “transition region”

still experience a low proportion of wetting storms. It is for these areas that it is difficult to produce sound estimates of the flood depth exceedance probability function. Figure 3.8 excludes areas where either all or no simulated storms produced measurable surge elevations (in other words, it only shows grid points with between 1 and 91 wetting observations) to offer a tighter understanding of where pseudo-surge values are intended to make improvements in the regression. The analysis is restricted to these areas for the remainder of the thesis.

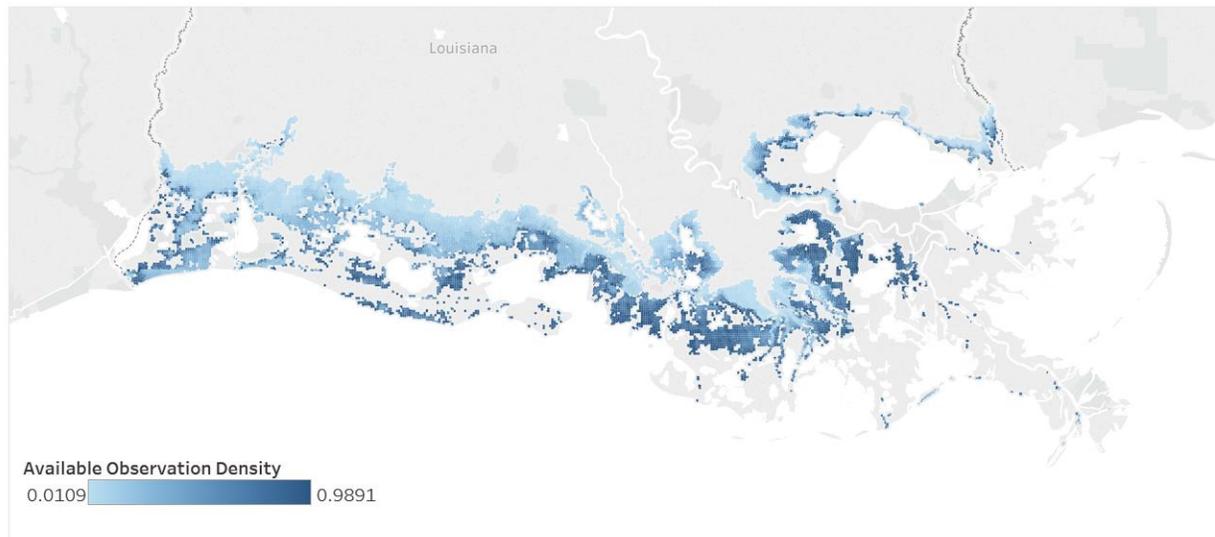


Fig. 3.8 Observation Density for Grid Points of Interest

Figure 3.9 gives a geographic representation of regression methods employed throughout the region of interest. This helps identify which sections of the coast use a lower priority regression in the response surface fitting procedure. Large portions of the region use locally weighted regression (1) which is implemented at the watershed geospatial unit of analysis. An identifiable pointwise full-form regression model (2) is rare. There are large regions which use the reduced-form pointwise regression (3) or the track-by-track regression (4) and very few locations which use a step function method (5).

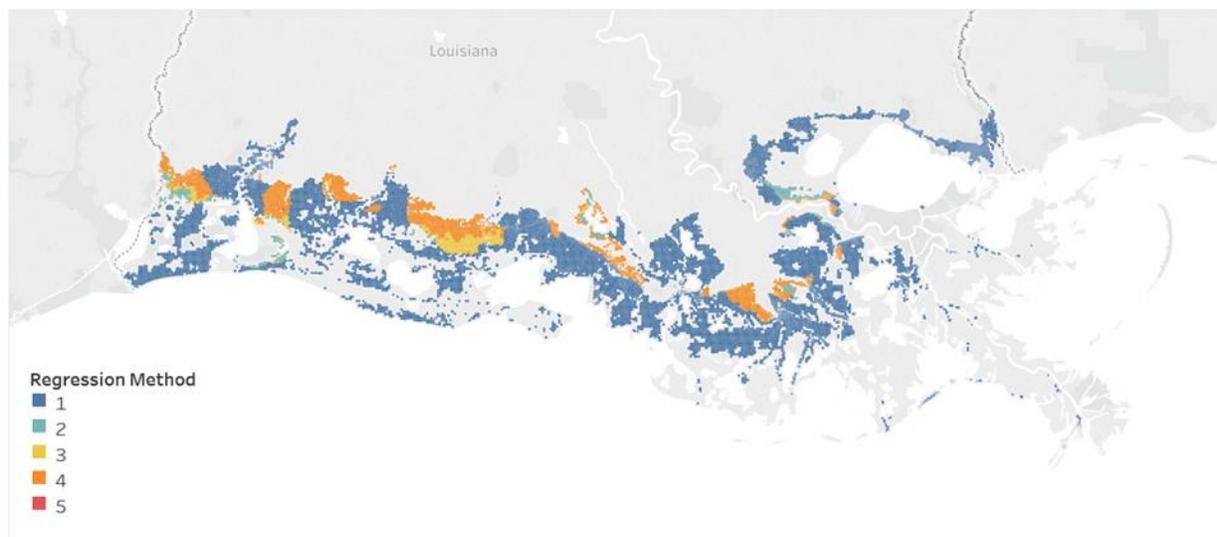


Fig. 3.9 Regression Method by Grid Point

It might be reasonable to expect pseudo-surge values to have a significant impact on exceedance curve estimates in areas which use regression methods 2-5. Pseudo-surge values may be less effective for watersheds which employ locally weighted regression because they typically already enjoy a large effective sample size. It is worth reiterating that a poor exceedance curve estimate can lead to poor estimates of expected annual damage, poor decisions about which flood mitigation projects should be implemented, and the possibility of high consequences when an intense storm event does occur. Additional insight regarding the CLARA 2017 response surface methodology is offered in APPENDIX C.

4. METHODS

The over-arching goal is to select pseudo-surge values for non-wetting storms in such a way as to minimize the error of the response surface estimates, thus reducing bias in flood exceedance curve estimates.

Up until this point, we have described different pieces of the CLARA statistical module which handles the data preprocessing portion of the annual exceedance probability function estimation. The diagram in 4.1 illustrates the CLARA 2017 production version workflow. The white boxes are tasks which are already well-defined and implemented and the black box is the task which is the focus of this thesis.

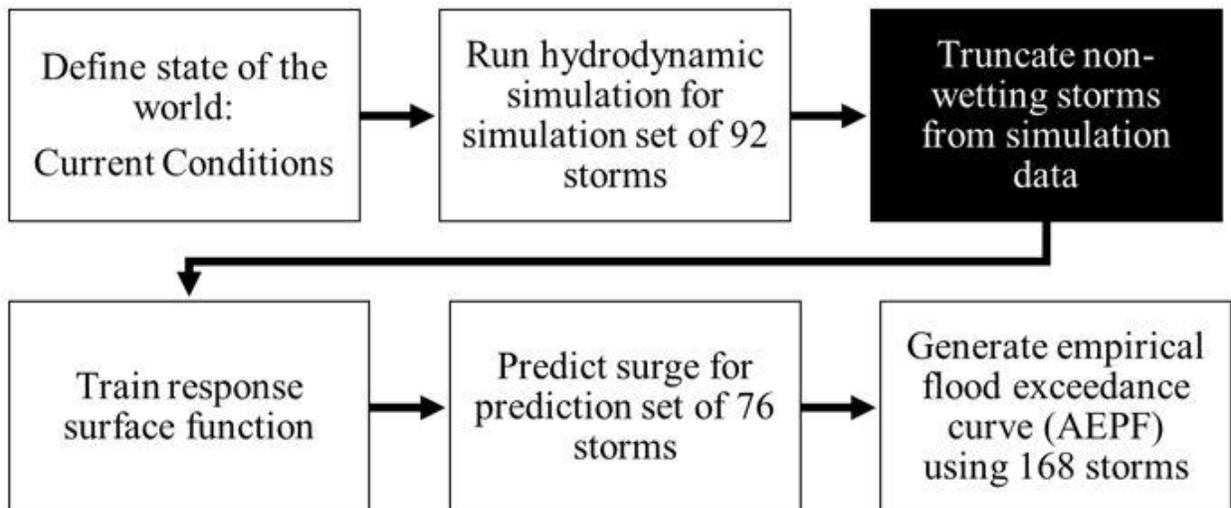


Fig. 4.1 Flood Depth Exceedance Probability Estimation Production Workflow

This thesis uses the current climatological conditions, coastal topography and bathymetry, as of 2015, to define the state of the world hereto referred to as “current conditions.” The ADCIRC+SWAN model simulations are available for the “gold standard” set of 446 storms in the JPM-OS suite of synthetic storms. These simulation outputs are interpreted as the baseline or ground-truth values with which the performance of response surface methodologies is measured, and best estimates of flood exceedance probability functions are generated.

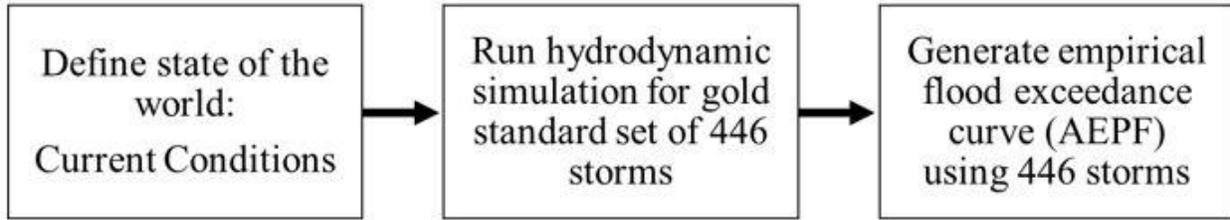


Fig. 4.2 Workflow to Generate Baseline True Surge and Flood Depth Exceedances

This allows us to quantify the value of information that could be gained by implementing pseudo-surge values.

Five methods for selecting pseudo-surge values are developed and/or investigated: 1) a topographic elevation rule, 2) a combined simulation and prediction set method, 3) a pointwise non-linear optimization, 4) a simultaneous watershed-level non-linear optimization, and 5) a heuristic approach to the simultaneous watershed-level method. These methods are developed to mirror the current response surface implementation in the CLARA model. Even though simulation data is available for all 446 storms, the investigation is focused on selecting optimal pseudo-surge values for the 92-storm subset used in the 2017 CMP analysis. The intent is to identify the impact of replacing the current “truncate non-wetting storms from simulation output” step in Fig. 4.1 with any one of the five pseudo-surge methods.

The resulting pseudo-surge values from each pseudo-surge selection method are used to train a corresponding response surface function using all 92 storms from the simulation set of storms regardless of their wetting status. For each new response surface function, surge estimates are made for all 76 storms in the prediction set resulting a total 168 surge estimates. Finally, separate annual exceedance probability function estimates are developed for each pseudo-surge selection method. The response surface functions and AEPF estimates are compared to the production version of the CLARA 2017 response surface methodology which truncates non-wetting storms from the response surface training procedure. This comparison is facilitated by the surge data collected across all 446 storms in the gold standard storm set.

4.1 Problem Description

Let A be the set of all locations and Ω be the set of all storms. For notational convenience, define $Z := A \times \Omega$. Consider the function $s: Z \rightarrow \mathbb{R}$ to be the true unknowable storm surge

function. Note that there exist some storms which do not generate surge at some locations, such that $s(z)$ is undefined. Often in this discussion we will refer to a partition of the domain Z into two sets, one of which refers to situations where surge is defined and another which refers to situations where surge is undefined. Let $\check{Z} = \{z \in Z : s(z) \text{ is defined}\}$ and $\tilde{Z} = \{z \in Z : s(z) \text{ is undefined}\}$.

Consider a simulation oracle which grants us the ability to observe $s(z)$ when it exists. Define a predictive storm surge function, $r(z|\boldsymbol{\beta})$, $r: Z \rightarrow \mathbb{R}$ parameterized by a real vector $\boldsymbol{\beta}$ and constrained to be a linear function of the form

$$r(z|\boldsymbol{\beta}) = \beta_0 + \beta_1 z_1 + \cdots + \beta_p z_p$$

where Z has dimension p , requiring us to estimate $p + 1$ coefficients. Our goal is to identify a “best-fit” vector of regression coefficients that minimizes the difference between $r(z)$ and $s(z)$.

In practice, it is not possible to know the value of $s(z)$ or even if it exists for all $z \in Z$, as this would require an infinite number of simulations. Instead, we use a discrete number of carefully chosen simulations to construct a “simulation” set M where $M \subset Z$. Note that there exists $z \in M$ such that $s(z)$ is undefined. Then partition M in the same manner as Z letting $\check{M} = \{z \in M : s(z) \text{ is defined}\}$ and $\tilde{M} = \{z \in M : s(z) \text{ is undefined}\}$. The regression error is defined to be

$$SSE_R := \sum_{z \in \check{M}} (s(z) - r(z|\boldsymbol{\beta}))^2 \quad (1)$$

If non-wetting storms are truncated, then the best fit vector $\boldsymbol{\beta}^*$ is obtained by minimizing SSE_R

$$\boldsymbol{\beta}^* := \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \sum_{z \in \check{M}} (s(z) - r(z|\boldsymbol{\beta}))^2. \quad (2)$$

This minimization includes only wetting storms, $z \in \check{M}$, and cannot include non-wetting storms, $z \in \tilde{M}$.

Once $\boldsymbol{\beta}^*$ is obtained, the function $r(z|\boldsymbol{\beta}^*)$ is used to make predictions of $s(z)$ for a discrete number of $z \in Z$ that form a “prediction” set $N \subset Z$. Note that the simulation prediction sets are disjoint ($M \cap N = \emptyset$), and that they do not partition the set of all location-storm pairs ($M \cup N \neq Z$). It is possible that there exists $z \in N$ such that $s(z)$ is undefined, but this is unknown because there are no simulations for $z \in N$. Nonetheless, partition N in the same

manner as Z and M letting $\check{N} = \{z \in N : s(z) \text{ is defined}\}$ and $\tilde{N} = \{z \in N : s(z) \text{ is undefined}\}$. Note that when Z , M , and N are partitioned, we only partition the storm set and not the location set.

Table 4.1 Storm-Location Set Definitions

Set	Description	Relationships
A	The set of all locations	-
Ω	The set of all storms	-
Z	The set of location storm pairs	$Z = A \times \Omega$
\check{Z}	The set of location storm pairs where the storm is wetting	$\check{Z} \subset Z$
\tilde{Z}	The set of location storm pairs where the storm is non-wetting	$\tilde{Z} \subset Z$ $\check{Z} \cup \tilde{Z} = Z$ $\check{Z} \cap \tilde{Z} = \emptyset$
M	The set of location storm pairs where the storm is in the simulation set of storms	$M \subset Z$
\check{M}	The set of location storm pairs where the storm is in the simulation set of storms and is wetting	$\check{M} \subset M$
\tilde{M}	The set of location storm pairs where the storm is in the simulation set of storms and is non-wetting	$\tilde{M} \subset M$ $\check{M} \cup \tilde{M} = M$ $\check{M} \cap \tilde{M} = \emptyset$
N	The set of location storm pairs where the storm is in the prediction set of storms	$N \subset Z$ $M \cap N = \emptyset$ $M \cup N \neq Z$
\check{N}	The set of location storm pairs where the storm is in the prediction set of storms and is wetting	$\check{N} \subset N$
\tilde{N}	The set of location storm pairs where the storm is in the prediction set of storms and is non-wetting	$\tilde{N} \subset N$ $\check{N} \cup \tilde{N} = N$ $\check{N} \cap \tilde{N} = \emptyset$

In a testing situation where $s(z)$ is known for $z \in N$, we define the prediction error to be

$$SSE_p := \sum_{z \in \check{N}} (s(z) - r(z|\boldsymbol{\beta}))^2 + \sum_{z \in \tilde{N}} \mathbb{I}(r(z|\boldsymbol{\beta}) > h_\alpha) (h_\alpha - r(z|\boldsymbol{\beta}))^2 \quad (3)$$

where h_α is the topographic elevation at location $\alpha \in A$, $\boldsymbol{\beta}$ is a vector of regression coefficients, and $\mathbb{I}(r(z|\boldsymbol{\beta}) > h_\alpha)$ is an indicator function for when the function $r(z|\boldsymbol{\beta})$ incorrectly predicts surge to exist for $z \in \tilde{N}$. The prediction error is a combination of estimate error for wetting storms and non-wetting storms in the prediction set. A measure of the total error is the sum of the regression error and the prediction error

$$SSE_T := \sum_{z \in \tilde{M} \cup \tilde{N}} (s(z) - r(z|\boldsymbol{\beta}))^2 + \sum_{z \in \tilde{M} \cup \tilde{N}} \mathbb{I}(r(z|\boldsymbol{\beta}) > h_\alpha) (h_\alpha - r(z|\boldsymbol{\beta}))^2. \quad (4)$$

A natural way to improve flood depth exceedance probability function estimates is to improve the similarity between the true storm surge function $s(z)$ and the predictive storm surge function $r(z)$. In the current response surface methodology, $\boldsymbol{\beta}^*$ is obtained using wetting storms in the simulation set, $s(z)$ for $z \in \tilde{M}$, which fails to take advantage of information from non-wetting storms in the simulation set, $z \in \tilde{N}$. The fact that $s(z)$ is undefined for $z \in \tilde{N}$ is itself valuable information.

At this point we introduce the concept of ‘‘pseudo-surge’’ $\tilde{s}(z)$ to be used for non-wetting storms in the simulation set, $z \in \tilde{N}$, with the requirement that a pseudo-surge value remain lower than the topographic elevation, $\tilde{s}(z) \leq h_\alpha$. Let these values be arranged in a vector, $\tilde{\mathbf{s}}$. The intent is to include pseudo-surge values through a rule of thumb which will improve the similarity between $s(z)$ and $r(z)$ through the regression coefficients. That is to use both $s(z)$ and $\tilde{s}(z)$ for $z \in M$ to determine the optimal regression coefficients

$$\tilde{\boldsymbol{\beta}}^* := \operatorname{argmin}_{\boldsymbol{\beta}} \left[\sum_{z \in \tilde{M}} (s(z) - r(z|\boldsymbol{\beta}))^2 + \sum_{z \in \tilde{N}} (\tilde{s}(z) - r(z|\boldsymbol{\beta}))^2 \right] \quad (5)$$

where $\tilde{\boldsymbol{\beta}}^*$ refers to optimal regression coefficients found using pseudo-surge values. We will use the ‘‘ $\tilde{\cdot}$ ’’ to denote instances where regression coefficients are generated under a pseudo-surge implementation.

In efforts to develop and evaluate this rule, we wish to find the pseudo-surge values which minimize the total error, SSE_T . For notational convenience consider the regression coefficients to be a function of $\tilde{\mathbf{s}}$, and the regression and prediction errors to be functions of $\tilde{\boldsymbol{\beta}}^*(\tilde{\mathbf{s}})$. Then the optimal pseudo-surge values solve the problem

$$\operatorname{argmin}_{\tilde{\mathbf{s}}} \left[\left(\min_{\tilde{\boldsymbol{\beta}}^*(\tilde{\mathbf{s}})} SSE_R \right) + SSE_P \right]. \quad (6)$$

This formulation maintains the fact that the regression coefficients of the prediction function are to be determined using data from the simulation set only.

4.2 Approaches to Identifying Pseudo-surge Values

Following are three approaches to identifying pseudo-surge values. In Section 4.2.1, we give information on topographic elevation replacement rule used in previous research. We then explore the possibility of simply extrapolating pseudo-surge values using all storms in the simulation and prediction sets in Section 4.2.2. This motivates Section 4.2.3; a non-linear optimization approach for selecting pseudo-surge values that is applied to multiple linear regression at the CLARA grid point unit of analysis. Finally, Sections 4.2.4 and 4.2.5 attempt to extend the non-linear optimization approach to the CPARLWR method at the watershed unit of analysis resulting in a straight-forward and heuristic algorithm method respectively.

4.2.1 Topographic Elevation Replacement Rule

The concept of replacing non-wetting surge observations with some other value has been implemented in a previous study of flood risk in Louisiana. The United States Army Corps of Engineers' (USACE) Louisiana Coastal Protection and Restoration (LACPR) Report's Hydraulics and Hydrology Technical Appendix (2009) alludes to a simple approach which replaces undefined observations with a value that is $3/10$ the topographic elevation. This is essentially an algorithm that selects pseudo-surge values at a location $\alpha \in A$ where non-wetting exists, $\tilde{M} \neq \emptyset$, such that $\tilde{s}(z) = (1/3)h_\alpha$. Consequently, the same pseudo-surge value is imputed for all non-wetting storms at the same location. Using such a ratio rule requires that an absolute zero is implied (Stevens, 1946). Here, issues arise when the zero value of sea level is defined by an arbitrary elevation datum. The topographic elevation rule fails when a location on the coast has a negative elevation relative to sea level. In such cases, taking $\tilde{s}(z) = (1/3)h_\alpha$ produces pseudo-surge values that indicate wetting. This is contrary to the results of the hydrodynamic simulation. While easy to implement, this method produces undesirable results, as shown in Chapter 5.

4.2.2 Combined Regression and Extrapolation

An intuitive approach is to generate pseudo-surge values using a response surface function that is trained with the combined simulation output from the 92-storm simulation set and the 76-storm prediction set. This time, optimal regression coefficients are defined through the solution to the problem

$$\beta_c^* := \operatorname{argmin}_{\beta} \sum_{z \in \tilde{M} \cup \tilde{N}} (s(z) - r(z|\beta))^2. \quad (7)$$

The difference here is that the summation is over defined observations of surge in the simulation and prediction sets, not only defined observations in the simulation set. Next, define pseudo-surge values as the response function surge estimate for observations where surge is undefined

$$s(z) = \tilde{s}(z) = r(z|\beta_c^*) \text{ for } z \in \tilde{M}.$$

Now, with pseudo-surge values for $z \in \tilde{M}$, optimal regression coefficients, $\tilde{\beta}^*$, are obtained using the entirety of the simulation data and equation (5).

Consider the dummy data $x = \{1.5, 3, 5, 7\}$ and $y = \{NA, 8, 16, 28\}$ where x is the independent variable and y is the response variable. The data is plotted in Figure 4.3.

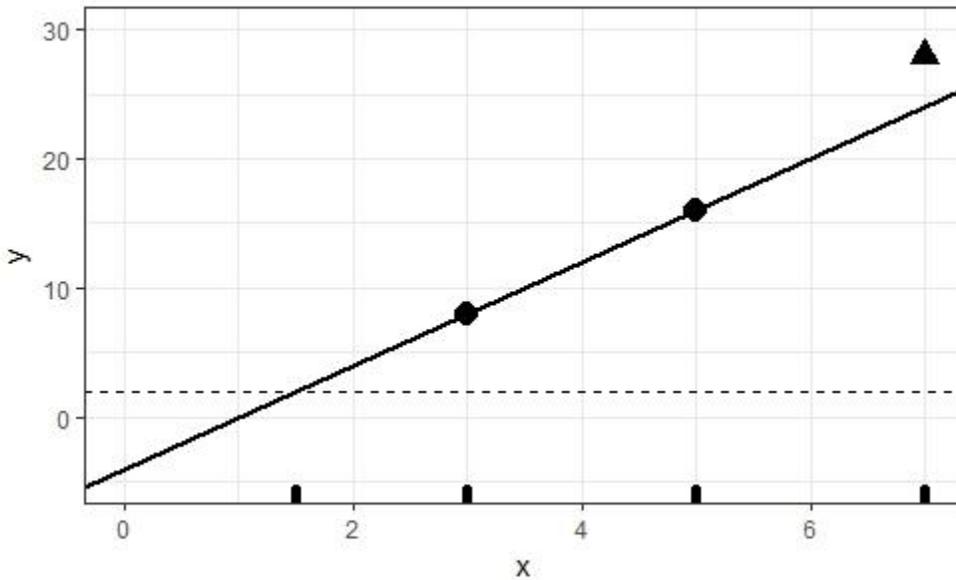


Fig. 4.3 Simple Extrapolation Example

Circular observations are those that are used to train a linear regression that predicts a value for the triangular observation. Note that one of the response values, y_1 , is undefined because its value fell below some threshold (analogous to topographic elevation) indicated by the horizontal dashed line at $y = 2$. This reduces the number of training observations to two. Fitting a linear model to the training data yields a regression line which is then used to predict a value for the triangular observation. The goal is to select a value for y_1 that reduces the regression error and the prediction error.

A candidate for y_1 must be less than or equal to the threshold value 2 to avoid misclassifying the observation as defined. Generating a pseudo-value from a linear regression fit using both simulation and prediction data yields $y_1 = -0.1167$. A new plot is shown in Figure 4.4 with the pseudo-value as a hollow circle. The new regression line, shown as dashed, is fit using the solid and hollow circular observations and has increased performance when predicting the triangular observation, though decreased performance for circular observations.

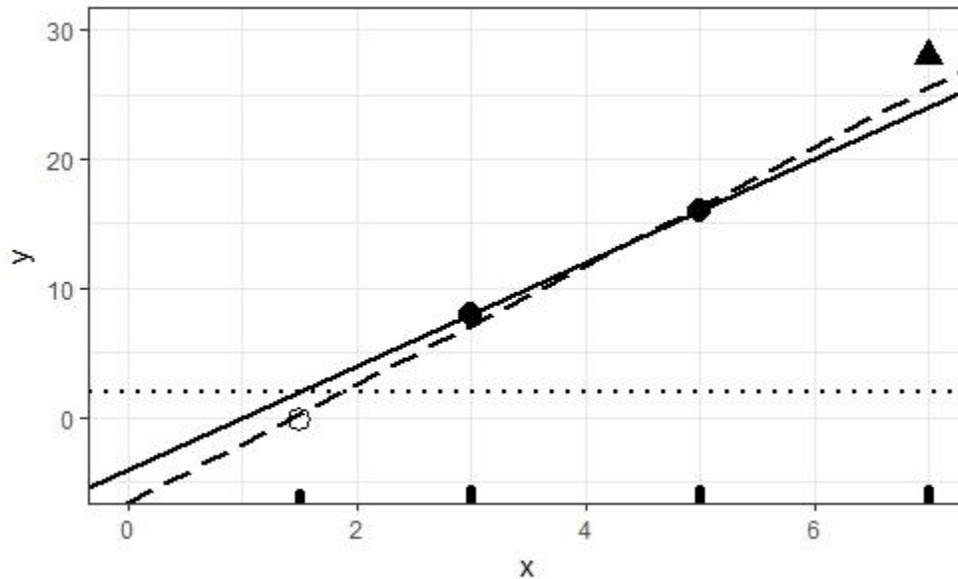


Fig. 4.4 Example Regression with Pseudo-value Included

While this procedure may provide a reasonable value for y_1 it may not be the optimal choice. That is, there may exist a different choice for y_1 for which the total error, SSE_T , for available observations is further minimized. To investigate this, consider the range of possible values for y_1 from -5 to 2 using a step-size of 0.05 .

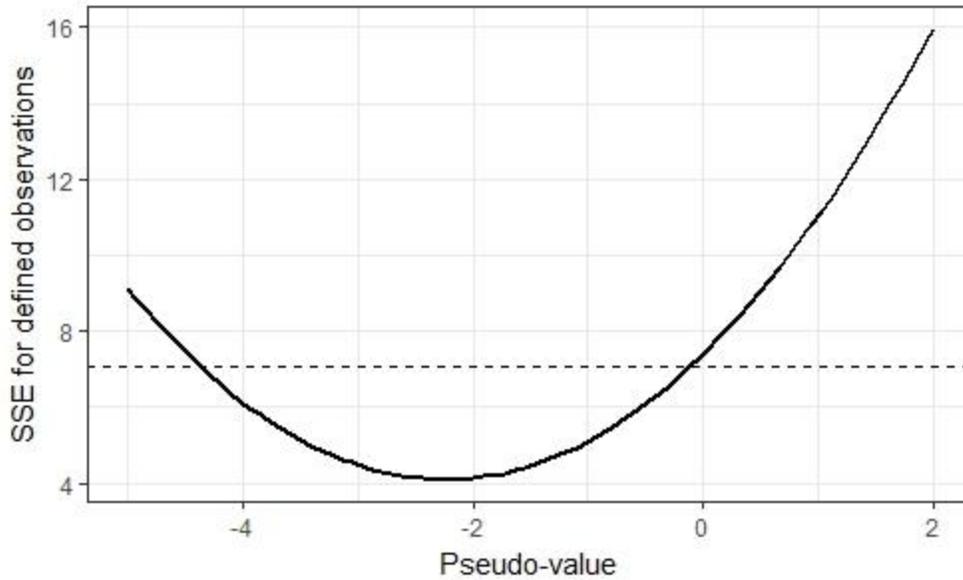


Fig. 4.5 SSE for Defined Observations as a Function of the Pseudo-value

Figure 4.5 shows the SSE_T for available observations when we let y_1 vary. At the originally selected value of $y_1 = -0.11\bar{6}$ (shown with the dashed line) the SSE_T is not minimized. The SSE_T is minimized at $y_1 = -2.25$.

There are two major issues with this approach to determine good choices of pseudo-surge values. First, for some locations $\alpha \in A$, there may not be a unique solution for optimal regression coefficients β_c^* even after including wetting observations from the prediction set, $z \in \tilde{N}$, when training the regression function. If it is not possible to obtain a unique solution for β^* , then it is not possible to generate pseudo-surge values for non-wetting storms in the simulation set, $z \in \tilde{M}$. This is likely to happen for locations with very few defined observations and it is those locations where the most benefit of a pseudo-surge implementation is desired. Second, such a procedure is shown to not minimize SSE_T .

This example is a special case because there is only one undefined observation. Often multiple observations from storms in the simulation set or the prediction set can be undefined. In practice, non-wetting storms in the simulation set are easily identified and their estimated surge is reset to be non-wetting regardless of the response surface function surge estimate. The same cannot be done for storms in the prediction set because we do not have the true simulated storm surge. The total error is redefined to be

$$SSE_T := \sum_{z \in \tilde{M} \cup \tilde{N}} (s(z) - r(z|\boldsymbol{\beta}))^2 + \sum_{z \in \tilde{N}} \mathbb{I}(r(z|\boldsymbol{\beta}) > h_\alpha) (h_\alpha - r(z|\boldsymbol{\beta}))^2 \quad (8)$$

with $\boldsymbol{\beta}$ being a vector of regression coefficients and the second summation is over non-wetting storms in the prediction set only. The next section will examine the convexity of this loss function and the application of non-linear optimization.

4.2.3 Pointwise Selection of Pseudo-surge Values

The example presented in Section 4.2.2 motivates a non-linear optimization approach to obtain the optimal pseudo-surge values for undefined observations from the ADCIRC+SWAN hydrodynamic simulation. In this section, each CLARA grid point is considered independently, corresponding to regression method (2), point-by-point full-form linear regression, outlined in Section 3.5 CLARA Response Surface Methodology. At second best, the response surface is fit via a “full-form” linear regression. This happens when the number of available grid points in a watershed falls below 50 or if the CPARLWR model is not identifiable due to lack of variation in covariates for storms which caused wetting. A grid point is considered available when at least one storm in the simulation set produces wetting. If there are no wetting storms in the simulation set, then the grid point is removed from the regression altogether (i.e. truncated). There is no selection of pseudo-surge values for grid points that have no wetting storms in the simulation or prediction sets because they would not improve the response surface or AEPF estimates.

At this stage it is more convenient to switch to vector notation when discussing a non-linear optimization approach. The set definitions from Table 4.1 still hold. Instead of referring to a surge value $s(z)$ as a function of a storm-location ordered pair, let \mathbf{s} represent a vector of storm surge values. Subscript notation is used when referring to surge values for all $z \in Z$ or some subset of Z . For example, \mathbf{s}_M and $\mathbf{s}_{\tilde{N}}$ refer to surge values for $z \in M$ and $z \in \tilde{N}$ respectively. Further, let \mathbf{X}_M be the covariate matrix for the simulation set of storm data and \mathbf{X}_N be the covariate matrix for the prediction set of storm data. A row in these matrices contain the “full-form” linear regression covariate information for one simulated surge observation. For example, the vector $\mathbf{s}_{\tilde{M}}$ is the vector of surge values which are defined for $z \in \tilde{M}$ and $\mathbf{X}_{\tilde{M}}$ is a matrix of corresponding covariate data. Given $\boldsymbol{\beta}^*$, surge estimates are then $\hat{\mathbf{s}}_{\tilde{M}} = \mathbf{X}_{\tilde{M}}\boldsymbol{\beta}^*$. The use of

subscript $\alpha \in A$ on terms is resisted in the interest of keeping the expressions as simple as possible, but the following discussion applies to a single location (i.e. CLARA grid point) $\alpha \in A$.

For a single independent location $\alpha \in A$, the L_2 error associated with fitting a linear regression to data $\mathbf{s}_{\bar{M}}$ and $\mathbf{X}_{\bar{M}}$ is given as

$$\begin{aligned} SSE_R &= \|\mathbf{s}_{\bar{M}} - \hat{\mathbf{s}}_{\bar{M}}\|_2^2 \\ &= \|\mathbf{s}_{\bar{M}} - \mathbf{X}_{\bar{M}}\boldsymbol{\beta}\|_2^2. \end{aligned}$$

The analytic solution for the vector $\boldsymbol{\beta}^*$ which minimizes SSE_R is

$$\begin{aligned} \boldsymbol{\beta}^* &= \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \|\mathbf{s}_{\bar{M}} - \mathbf{X}_{\bar{M}}\boldsymbol{\beta}\|_2^2 \\ &= (\mathbf{X}_{\bar{M}}^T \mathbf{X}_{\bar{M}})^{-1} \mathbf{X}_{\bar{M}}^T \mathbf{s}_{\bar{M}}. \end{aligned}$$

The goal is to construct a vector of optimal coefficients, $\tilde{\boldsymbol{\beta}}^*$, using the entirety of the simulation set data \mathbf{s}_M and \mathbf{X}_M . Let $\mathbf{W} = (\mathbf{X}_M^T \mathbf{X}_M)^{-1} \mathbf{X}_M^T$ and let $\tilde{\mathbf{s}}$ be a vector of decision variables representing an ideal form of \mathbf{s}_M where all entries are defined either with an original surge value or with a pseudo-surge value. The analytic solution for the vector $\tilde{\boldsymbol{\beta}}^*$ is

$$\begin{aligned} \tilde{\boldsymbol{\beta}}^* &= (\mathbf{X}_M^T \mathbf{X}_M)^{-1} \mathbf{X}_M^T \tilde{\mathbf{s}} \\ &= \mathbf{W}\tilde{\mathbf{s}}. \end{aligned}$$

The regression error is then

$$SSE_R = \|\mathbf{s}_{\bar{M}} - \mathbf{X}_{\bar{M}}\mathbf{W}\tilde{\mathbf{s}}\|_2^2.$$

The prediction error is

$$SSE_P = \|\mathbf{s}_{\bar{N}} - \mathbf{X}_{\bar{N}}\mathbf{W}\tilde{\mathbf{s}}\|_2^2 + \|\mathbb{I}(\mathbf{X}_{\bar{N}}\mathbf{W}\tilde{\mathbf{s}} > h_\alpha)(h_\alpha - \mathbf{X}_{\bar{N}}\mathbf{W}\tilde{\mathbf{s}})\|_2^2$$

where $\mathbb{I}(\mathbf{X}_{\bar{N}}\mathbf{W}\tilde{\mathbf{s}} > h_\alpha)$ is an indicator vector with elements equal to 1 when the corresponding surge estimate is above the topographic elevation and 0 otherwise. Note that regression error is still measured for wetting simulation set storms only. The first term of the prediction error is for wetting storms, while the second term measures error for non-wetting storms. The vectors $\mathbf{s}_{\bar{M}}$ and $\mathbf{s}_{\bar{N}}$ and matrices $\mathbf{X}_{\bar{M}}$ and $\mathbf{X}_{\bar{N}}$ are concatenated to produce $\mathbf{s}_{\bar{M}\cup\bar{N}}$ and $\mathbf{X}_{\bar{M}\cup\bar{N}}$ respectively. The total error in vector notation is given as

$$SSE_T = \|\mathbf{s}_{\bar{M}\cup\bar{N}} - \mathbf{X}_{\bar{M}\cup\bar{N}}\mathbf{W}\tilde{\mathbf{s}}\|_2^2 + \|\mathbb{I}(\mathbf{X}_{\bar{N}}\mathbf{W}\tilde{\mathbf{s}} > h_\alpha)(h_\alpha - \mathbf{X}_{\bar{N}}\mathbf{W}\tilde{\mathbf{s}})\|_2^2.$$

Choosing pseudo-surge values for entries of $\tilde{\mathbf{s}}$ is then equivalent to solving the problem

$$\begin{aligned} \text{P1: } \tilde{\mathbf{s}}^* &= \underset{\tilde{\mathbf{s}}}{\text{argmin}} \|\mathbf{s}_{\tilde{M} \cup \tilde{N}} - \mathbf{X}_{\tilde{M} \cup \tilde{N}} \mathbf{W} \tilde{\mathbf{s}}\|_2^2 + \|\mathbb{I}(\mathbf{X}_{\tilde{N}} \mathbf{W} \tilde{\mathbf{s}} > h_\alpha)(h_\alpha - \mathbf{X}_{\tilde{N}} \mathbf{W} \tilde{\mathbf{s}})\|_2^2 \\ &\text{s.t.} \\ &\tilde{\mathbf{s}} \leq h_\alpha \text{ for } z \in \tilde{M} \\ &\tilde{\mathbf{s}} = s_M \text{ for } z \in \tilde{M} \\ &\mathbf{W}_{c_p} \tilde{\mathbf{s}} \leq 0 \\ &\mathbf{W}_{r_{max}} \tilde{\mathbf{s}} \geq 0 \end{aligned}$$

where \mathbf{W}_{c_p} and $\mathbf{W}_{r_{max}}$ are the rows of \mathbf{W} which correspond to the central pressure and radius of maximum windspeed covariates.

The first two constraints indicate that for all non-wetting storms (undefined surge), the pseudo-surge value must be less than the topographic elevation, h_α , and that for all wetting storms (defined surge), the corresponding decision variable must equal the original surge value from s_M . The last two constraints require that the estimated surge must have a negative relationship with central pressure and a positive relationship with the radius of maximum windspeed.

Solving this minimization problem requires the implementation of a non-linear optimization algorithm such as sequential quadratic programming (SQP) or the interior-point method (Nordecal, 2006). The regression and prediction error term in this non-linear optimization invokes the L_2 vector norm which is known to be convex (Krantz, 2015). It is left to show that the error for non-wetting storms in the prediction set, $\|\mathbb{I}(\mathbf{X}_{\tilde{N}} \mathbf{W} \tilde{\mathbf{s}} > h_\alpha)(h_\alpha - \mathbf{X}_{\tilde{N}} \mathbf{W} \tilde{\mathbf{s}})\|_2^2$, is also convex.

The second term involves a vector of indicator functions. A non-wetting storm in the prediction set is correctly classified as non-wetting when the surge estimate is below the topographic elevation and error is set to 0. However, squared error accumulates starting from 0 when the surge estimate is above the topographic elevation. For simplicity in explanation, assume that there is only one non-wetting storm in the prediction set for which to measure prediction error. Then the prediction error for this storm can be defined as a piecewise function

$$f(\hat{s}) = \begin{cases} 0, & \hat{s} \leq h_\alpha \\ (h_\alpha - \hat{s})^2, & \hat{s} > h_\alpha \end{cases}$$

where \hat{s} represents the surge estimate for the non-wetting storm.

The convexity of $f(\hat{s})$ is shown. It is required to show that for any $\hat{s}_1, \hat{s}_2 \in \mathbb{R}$ and any $\lambda \in [0,1]$, it holds that

$$f((1-\lambda)\hat{s}_1 + \lambda\hat{s}_2) \leq (1-\lambda)f(\hat{s}_1) + \lambda f(\hat{s}_2).$$

Let $\hat{s}_1 \leq \hat{s}_2$. In the case that $\hat{s}_1, \hat{s}_2 \leq h_\alpha$, then f is a constant function and is convex (though not strictly convex). In the case that $\hat{s}_1, \hat{s}_2 > h_\alpha$, then f is parabolic with a positive coefficient and is strictly convex. For the final case, let $\hat{s}_1 \leq h_\alpha$ and $\hat{s}_2 > h_\alpha$. For $\lambda \in [0,1]$ such that $((1-\lambda)\hat{s}_1 + \lambda\hat{s}_2) \leq h_\alpha$, then

$$\begin{aligned} f((1-\lambda)\hat{s}_1 + \lambda\hat{s}_2) &= 0 \\ &\leq (1-\lambda)f(\hat{s}_1) + \lambda f(\hat{s}_2) \\ &= 0 + \lambda(h_\alpha - \hat{s}_2)^2 \\ &= \lambda(h_\alpha - \hat{s}_2)^2. \end{aligned}$$

For $\lambda \in [0,1]$ such that $((1-\lambda)\hat{s}_1 + \lambda\hat{s}_2) > h_\alpha$, then

$$\begin{aligned} f((1-\lambda)\hat{s}_1 + \lambda\hat{s}_2) &= (h_\alpha - (1-\lambda)\hat{s}_1 - \lambda\hat{s}_2)^2 \\ &\leq (1-\lambda)f(\hat{s}_1) + \lambda f(\hat{s}_2) \\ &= 0 + \lambda(h_\alpha - \hat{s}_2)^2 \\ &= \lambda(h_\alpha - \hat{s}_2)^2. \end{aligned}$$

Thus, f is a convex function.

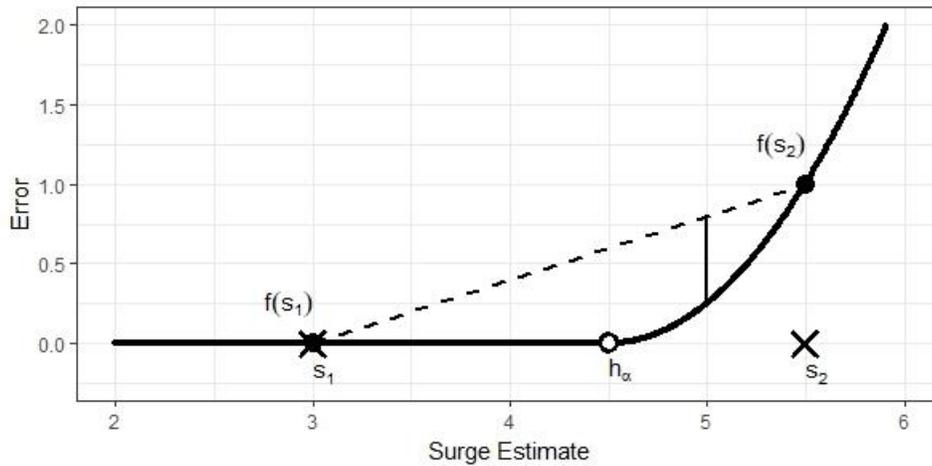


Fig. 4.6 Convexity of $f(\hat{s})$

Figure 4.6 offers a visual representation. The argument can be extended to loss functions with vector inputs such as those present in P1. The function f can be considered one-sided L_2

loss and is not necessarily foreign to the field of classification problems. Other loss functions for classification exist such as indicator, logistic, and hinge loss (Rosasco et al., 2003). One benefit of one-sided square error loss is that it severely penalizes gross misclassification of non-wetting storms. As a sum of convex functions, the objective function is convex. Additionally, all constraints are linear. Thus, the problem is amenable to standard non-linear convex optimization algorithms.

This is a pointwise approach that is independently applied to each CLARA grid point sequentially. This corresponds to method 2, point-by-point full-form linear regression, of the CLARA response surface methodology and at best can allow a point-specific full-form regression model.

4.2.4 Simultaneous Selection of Pseudo-surge Values

The preceding section developed a method to select pseudo-surge values in a pointwise manner. However, the regression method with the highest priority in the CLARA response surface is conditionally parametric locally weighted regression (CPARLWR) (Fischbach et al., 2017). It is desirable to extend the pointwise method to one which can simultaneously select pseudo-surge values for all undefined observations in a watershed. Since the regression is applied to a geographic setting, locally weighted regression is often termed geographically weighted regression. A response variable at a target location may be influenced by observations in some local vicinity of the target location. Locally weighted regression accounts for this fact by appropriately weighting local information and including it in the regression at the target location.

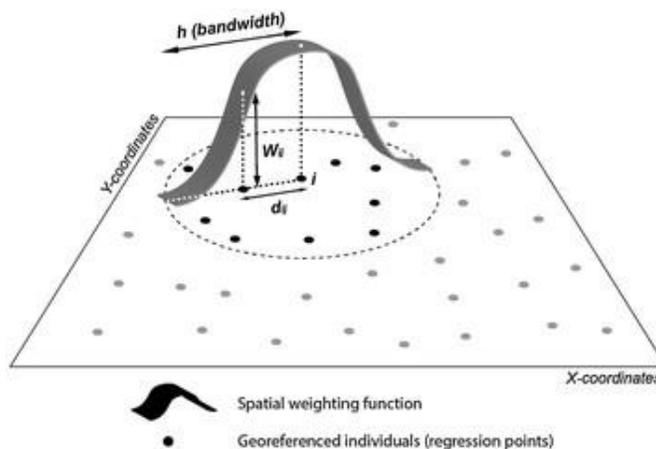


Fig. 4.7 2-D Locally Weighted Regression Example (Feuillet et al., 2015)

Setting aside the previous notation for a moment, Figure 4.7 shows a typical locally weighted regression scheme. For a target point j , a spatial weighting function with bandwidth h assigns a weight w_j to location i according to the distance d_i from i to j . The grey points are those which are considered non-local to the target point j because they lie beyond some maximum distance h from j . A regression with weighted observations is carried out much the same way as multiple linear regression except the coefficient estimates are calculated by

$$\hat{\beta}_j = (\mathbf{X}^T \mathbf{W}_j \mathbf{X})^{-1} \mathbf{X}^T \mathbf{W}_j \mathbf{y},$$

where \mathbf{W}_j is a diagonal matrix of the weights assigned to each realization of covariate values stored in a row of \mathbf{X} . These regression coefficients are defined for location j , the target location. The algorithm then moves to a new target location and executes the same procedure to develop the location-specific vector of coefficient estimates. The CLARA response surface methodology invokes CPARLWR if there are enough grid points available within a watershed. The CLARA implementation uses a tri-cube spatial weighting function to define response weights and selects the bandwidth parameter in a general cross-validation procedure (Fischbach et al., 2017).

For watersheds in which CPARLWR is invoked, the goal is to simultaneously select pseudo-surge values because the CLARA grid points are no longer considered to be independent of each other.

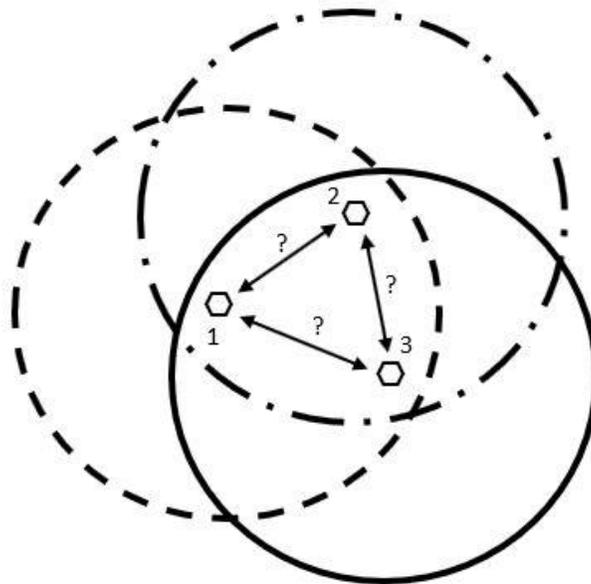


Fig. 4.8 Unknown Impacts of Pointwise Pseudo-surge on LWR

Consider the orientation of three georeferenced points within a region in Figure 4.8, where the corresponding dashed circles define the bandwidth of a local regression scheme. Each point is within the bandwidth of every other point so the regression at each point will include weighted information from all points. If point 3 were to have some number of undefined observations, then pointwise selection of pseudo-surge may produce values that negatively impact the weighted regressions at points 1 and 2. For this reason, the goal is then to select pseudo-surge values simultaneously for all unavailable observations in a region.

The CPARLWR method is applied at the watershed geospatial unit of analysis because each watershed is a region of independent hydrodynamic activity. The CLARA grid points within a watershed region are the target locations of the geographically weighted regression and are not considered independent of each other. The collection of k watersheds is a partition of the set of all locations A . Define an arbitrary subregion $K \subset A$ as one such watershed containing a discrete number of n locations (CLARA grid points). The goal is to minimize the aggregate regression and prediction error across all locations $\alpha \in K$,

$$\begin{aligned} SSE_{T,K} &= SSE_{R,K} + SSE_{P,K} \\ &= \sum_{\alpha \in K} SSE_{R,\alpha} + \sum_{\alpha \in K} SSE_{P,\alpha} \end{aligned}$$

Where $SSE_{T,\alpha}$, $SSE_{R,\alpha}$, and $SSE_{P,\alpha}$ are the total, regression, and prediction sum of squared errors respectively at location $\alpha \in K$.

The approach is analogous to the pointwise selection of pseudo-surge with a few tweaks. Previously, \mathbf{s} was defined to be the vector of storm surge values for a single location $\alpha \in A$. Now, let \mathbf{s} be a much larger vector of storm surge values for all $\alpha \in K$. Likewise, let \mathbf{X} be a matrix of covariate data corresponding to entries in \mathbf{s} . Let \mathbf{D}_α be defined as a diagonal matrix of weights for the weighted regression at location $\alpha \in K$. Diagonal entries of \mathbf{D}_α are zero when corresponding rows of \mathbf{X} are considered non-local to the target location α . This allows use of a single vector and matrix for all surge observations and covariate data in a subregion.

The analytic solution for regression coefficients, $\boldsymbol{\beta}_\alpha^*$, that minimize the weighted regression error for observations local to the location α is

$$\begin{aligned} \boldsymbol{\beta}_\alpha^* &= \underset{\boldsymbol{\beta}}{\operatorname{argmin}} SSE_{R,\alpha} \\ &= \left(\mathbf{X}_{\bar{M}}^T \mathbf{D}_\alpha \mathbf{X}_{\bar{M}} \right)^{-1} \mathbf{X}_{\bar{M}}^T \mathbf{D}_\alpha \mathbf{s}_{\bar{M}}. \end{aligned}$$

Just as before, the goal is to construct the vector of optimal coefficients, $\tilde{\boldsymbol{\beta}}_\alpha^*$, using the entirety of the data \mathbf{s}_M and \mathbf{X}_M . Let $\mathbf{W}_\alpha = (\mathbf{X}_M^T \mathbf{D}_\alpha \mathbf{X}_M)^{-1} \mathbf{X}_M^T \mathbf{D}_\alpha$ and let $\tilde{\boldsymbol{\beta}}_\alpha$ be a vector of decision variables representing an ideal form of \mathbf{s}_M where all entries are defined either with an original surge value or with a pseudo-surge value. The analytic solution for the vector $\tilde{\boldsymbol{\beta}}_\alpha^*$ is

$$\begin{aligned}\tilde{\boldsymbol{\beta}}_\alpha^* &= (\mathbf{X}_M^T \mathbf{D}_\alpha \mathbf{X}_M)^{-1} \mathbf{X}_M^T \mathbf{D}_\alpha \tilde{\boldsymbol{\beta}}_\alpha \\ &= \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha.\end{aligned}$$

Then the regression error across all locations in the region is

$$SSE_{R,K} = \sum_{\alpha \in K} \|\mathbf{s}_{\tilde{M},\alpha} - \mathbf{X}_{\tilde{M},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha\|_2^2.$$

The total error for the region in the same fashion as before,

$$SSE_{T,K} = \sum_{\alpha \in K} \left(\|\mathbf{s}_{\tilde{M} \cup \tilde{N},\alpha} - \mathbf{X}_{\tilde{M} \cup \tilde{N},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha\|_2^2 + \|\mathbb{I}(\mathbf{X}_{\tilde{N},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha > h_\alpha)(h_\alpha - \mathbf{X}_{\tilde{N},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha)\|_2^2 \right).$$

Choosing pseudo-surge values for entries of $\tilde{\boldsymbol{\beta}}_\alpha$ is then equivalent to solving the problem

$$\begin{aligned}\text{P2: } \tilde{\boldsymbol{\beta}}_\alpha^* &= \underset{\tilde{\boldsymbol{\beta}}_\alpha}{\text{argmin}} \sum_{\alpha \in K} \left(\|\mathbf{s}_{\tilde{M} \cup \tilde{N},\alpha} - \mathbf{X}_{\tilde{M} \cup \tilde{N},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha\|_2^2 + \|\mathbb{I}(\mathbf{X}_{\tilde{N},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha > h_\alpha)(h_\alpha - \mathbf{X}_{\tilde{N},\alpha} \mathbf{W}_\alpha \tilde{\boldsymbol{\beta}}_\alpha)\|_2^2 \right) \\ &\text{s.t.} \\ &\tilde{\boldsymbol{\beta}}_{\tilde{M}} \leq h_\alpha \\ &\tilde{\boldsymbol{\beta}}_{\tilde{M}} = \mathbf{s}_M \\ &\mathbf{W}_{c_p,\alpha} \tilde{\boldsymbol{\beta}}_\alpha \leq 0 \text{ for } \alpha \in K \\ &\mathbf{W}_{r_{max},\alpha} \tilde{\boldsymbol{\beta}}_\alpha \geq 0 \text{ for } \alpha \in K\end{aligned}$$

where the constraints match those from the problem presented in Section 4.2.3 but are also location dependent.

In this method, pseudo-surge values are chosen for all undefined surge observations in a watershed at the same time. The objective function is convex because it is the sum of convex functions (Krantz, 2015). Obtaining the solution is achievable in theory, but in practice the length of \mathbf{s} and number of locations in K can make the optimization prohibitively expensive given available resources. Over 50 watersheds are fit using CPARLWR in the CLARA model. Table 4.2 gives a summary of just four of these. The number of total observations is the required length of the decision vector $\tilde{\boldsymbol{\beta}}_\alpha$ for the simultaneous selection of pseudo-surge values. Alternate formulations of the same problem may allow $\tilde{\boldsymbol{\beta}}_\alpha$ to be reduced to at least the number of undefined

observations. However, in the worst case among these selected watersheds, this would still require a decision vector $O(10^5)$ in length.

Table 4.2 Summary Characteristics for Select Watersheds

ID	Number of points	Number of local points	Defined Obs	Undefined Obs	Total Obs	Observation Density
77	1,418	36	25,850	104,606	130,456	0.20
86	238	24	5,071	16,825	21,896	0.23
180	5,135	26	113,528	358,892	472,420	0.24
64	3,026	61	100,408	177,984	278,392	0.36

A full table of summary characteristics of all watersheds is provided in APPENDIX A.

4.2.5 Ad Hoc Approach to Simultaneous Selection of Pseudo-Surge Values

Solutions for the simultaneous selection of pseudo-surge values for a watershed are computationally costly to obtain given available computing resources. Instead, the following section presents an ad hoc approach to the same problem, though potentially at the cost of producing sub-optimal candidate pseudo-surge values.

Finding the optimal solution for the simultaneous selection of pseudo-surge is difficult because the watersheds are composed of many georeferenced points. This problem could be mitigated by partitioning a watershed into many smaller regions where selecting an optimal pseudo-surge vector is less computationally costly. However, these sub-watersheds do not have independent hydrology which could cause the regression and prediction error to increase along internal boundaries. Furthermore, determining the “best” partition of a watershed is itself a challenging question. Instead, a convenient characteristic of the locally weighted regression is used to define a set of overlapping sub-regions and ensure that overlaps are considered when choosing a pseudo-surge value for a specific undefined observation.

For each grid point, a subregion is defined using the bandwidth, h , from the CPARLWR scheme. In this way, a set of overlapping subregions are created each with a CLARA grid point at its center. The intent is to improve the fit of a single weighted linear regression at the center grid point of each subregion using weighted observations from all other points in the subregion where the weights are determined by the same tri-cube kernel weighting function. The pointwise selection of pseudo-surge is applied to the subregion center point only but is still allowed to assign pseudo-surge values for any observations in the subregion that are undefined. Both the

pseudo-surge value and the regression weight is recorded for each of these undefined observations. The overlapping nature of the subregions requires that many pseudo-surge values will be separately selected, with corresponding weights, for the same undefined observation. The final pseudo-surge value for an undefined observation is then chosen to be the weighted average of all separately selected pseudo-surge values.

An example of this process is presented in Figure 4.9 with corresponding Table 4.3. Here there are five georeferenced points and point 3 has a single undefined observation that we wish to replace with a pseudo-surge value. The procedure iterates through the weighted regression at each target point, selecting the optimal pseudo-surge value for the undefined observation at point 3 and records the corresponding weight. In this example, four different pseudo-surge values (with weights) are selected for the same undefined observation. The final pseudo-surge value is the weighted average of these four values.

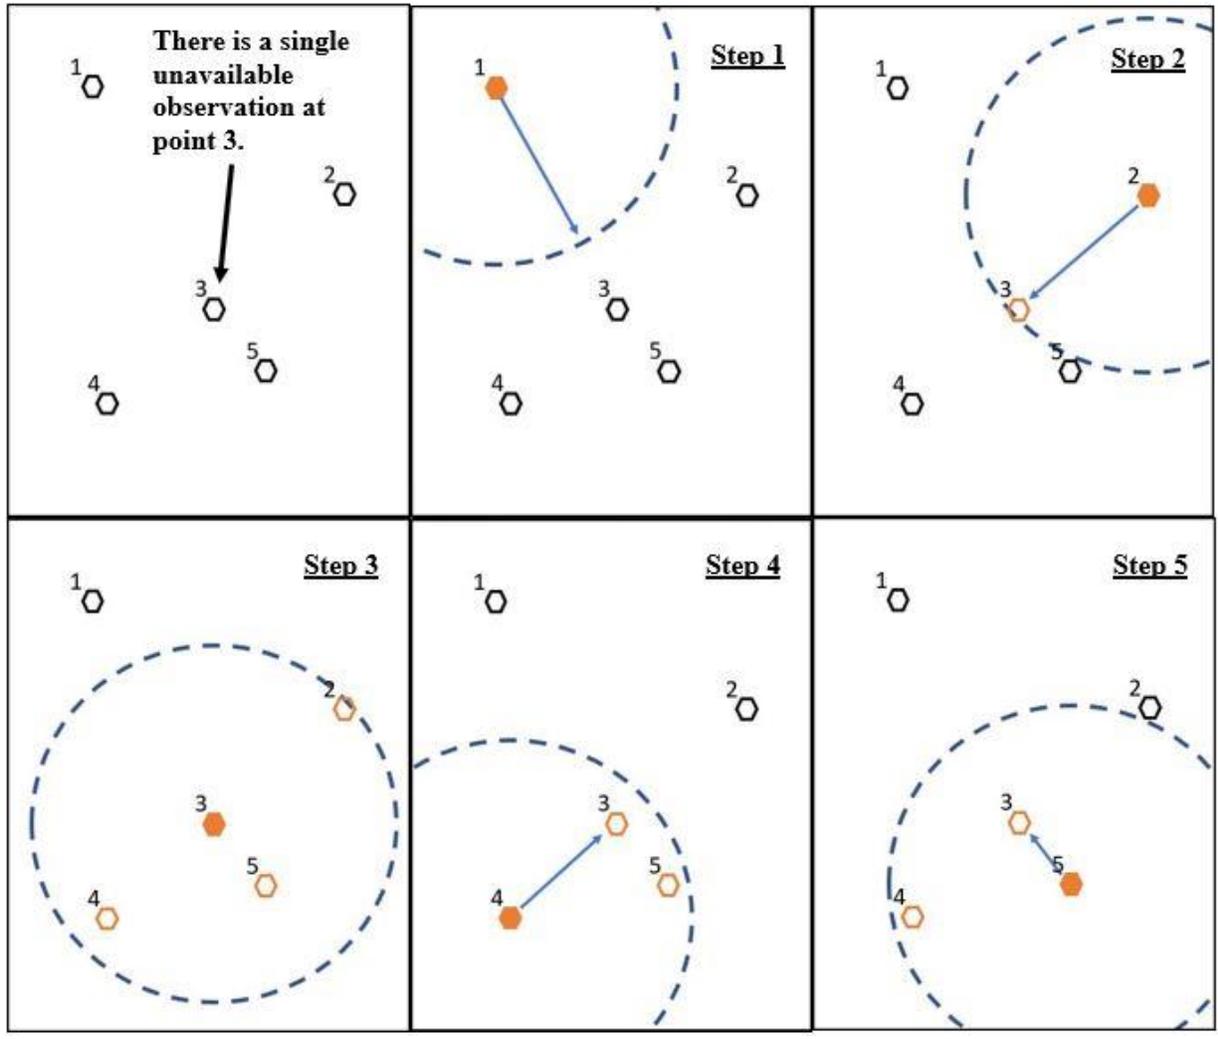


Fig. 4.9 Ad Hoc Simultaneous Selection of Pseudo-surge Visualization

Table 4.3 Ad Hoc Simultaneous Selection of Pseudo-surge Tabulation

Target Point (Step)	Pseudo-surge (ft)	Weight
1	Non-local	-
2	-5.00	0.02
3	-7.00	1.00
4	-2.00	0.20
5	-9.00	0.50
Pseudo-surge (ft)		-6.98

The procedure is outlined in the following algorithm:

Algorithm: Ad Hoc Selection of Pseudo-Surge Values

Input: Surge response vector \mathbf{s} , covariate data matrix \mathbf{X} , matrices \mathbf{W}_α and weight matrices \mathbf{D}_α
 $\alpha \in K$

Output: Pseudo-surge values $\tilde{s}(z)$

For each $\alpha \in K$

$$\text{Solve: } \underset{\tilde{\mathbf{s}}}{\text{argmin}} \left\| \mathbf{s}_{\tilde{M} \cup \tilde{N}, \alpha} - \mathbf{X}_{\tilde{M} \cup \tilde{N}, \alpha} \mathbf{W}_\alpha \tilde{\mathbf{s}} \right\|_2^2 + \left\| \mathbb{I}(\mathbf{X}_{\tilde{N}, \alpha} \mathbf{W}_\alpha \tilde{\mathbf{s}} > h_\alpha) (h_\alpha - \mathbf{X}_{\tilde{N}, \alpha} \mathbf{W}_\alpha \tilde{\mathbf{s}}) \right\|_2^2$$

s.t.

$$\tilde{\mathbf{s}}_{\tilde{M}} \leq h_\alpha$$

$$\tilde{\mathbf{s}}_{\tilde{M}} = s_M$$

$$\mathbf{W}_{c_p, \alpha} \tilde{\mathbf{s}} \leq 0 \text{ for } \alpha \in K$$

$$\mathbf{W}_{R_{max}, \alpha} \tilde{\mathbf{s}} \geq 0 \text{ for } \alpha \in K$$

Next α

For each $z \in \tilde{M} \times K$

Compile all pseudo-surge values assigned to $s(z)$ in vector \mathbf{y}_z

Compile weights assigned to $s(z)$ in vector \mathbf{d}_z

$$\text{Set } \tilde{s}(z) = \frac{\mathbf{y}_z \cdot \mathbf{d}_z}{\text{sum}\{\mathbf{d}_z\}}$$

Next z

4.3 Implementation

The methods described above were implemented using a combination of MATLAB and R programs. The production version of the CLARA 2017 response surface methodology execution and flood exceedance curve generation is handled within the R statistical software. Data preprocessing for the optimization procedures is also completed in R in order to lighten the computational requirements of MATLAB. Specifically, the partitioning of datasets, definition of an initial solution, calculation of weight and coefficient matrices (\mathbf{D}_α and \mathbf{W}) are all handled in R. Any matrix multiplication tasks that are not directly related to the vector of optimal pseudo-surge values are precomputed in this manner and packaged in MATLAB data files. The optimization which selects pseudo-surge values is executed MATLAB using the “fmincon” or “function minimization with constraints” routine. The results in the following section were obtained using the SQP algorithm invoked by the fmincon routine (Nelder, 2006). Upper and lower bounds, constraints, and objective function are prepared in MATLAB using the data

preprocessed in R. To cope with limited memory requirements, CLARA grid points and watersheds are separated across multiple data files which are operated on individually by MATLAB and then aggregated in post-optimization processing in R. Finally, R is used again to train separate CLARA response surfaces under each pseudo-surge implementation and the results go on to generate new flood exceedance curves.

The workflow required a seamless transition between R and MATLAB to operate effectively and was facilitated by the ability for R to write data in the MATLAB data file format. The optimization routine may benefit from reimplementing in a language designed to better handle large scale problems. Such a reimplementing may allow to use of the simultaneous selection of pseudo-surge rather than the proposed heuristic method in Section 4.2.5. However, because the CLARA model relies on R to train the response surface and generate exceedance curves, a workflow would need to be designed to handle the transition of data from one environment to the other.

5. RESULTS & DISCUSSION

The methodology section discussed five different methods to be applied to the problem of undefined hydrodynamic simulation output. The impact that these methods have on the response surface performance and the annual flood depth exceedance probabilities are now compared using simulation results from the “gold standard” set of 446 storms under current conditions as the baseline. This data is treated as the true surge experienced at each CLARA grid point for each of the 446 storms. These true surge values are used to measure the regression and prediction error of the response surface with or without a pseudo-surge implementation. Additionally, baseline annual flood depth exceedance curves are generated with the gold standard storms. These baseline exceedance curves are used to measure the absolute bias of exceedance curves generated with any other method.

The results section is broken into two parts. First, Section 5.2 shows performance of the response surface methodology measured in root mean squared error (RMSE) and storm classification rates with and without pseudo-surge implementations. Second, Section 5.3 shows impacts on flood depth exceedance probability estimates with and without pseudo-surge implementations. It is important to note that the response surface methodology uses 168 storms (92 simulated storms and 76 predicted storms) of 446 gold standard storm set which causes an underlying bias in the results regardless of pseudo-surge implementation. Further, the state-of-the-art response surface methodology used in the 2017 CLARA model is referred to as the “production” method. The production method truncates non-wetting storms from the response surface fitting procedure.

This thesis investigates the impact of implementing pseudo-surge values, thus focus is restricted to CLARA grid points where surge is produced for between 1 and 91 storms in the simulation set. Additionally, we do not use the ad hoc simultaneous selection of pseudo-surge method to determine pseudo-surge values for watersheds which were not deemed “large enough” to invoke the CPARLWR method. In these cases, no comparison can be made between a pointwise selection method and a simultaneous selection method.

A final note is the difficulty in summarizing and visualizing results. There is a danger of overlooking regions which suffer from high regression error or poor exceedance probability estimates when using metrics aggregated across the full study region. Additionally, each CLARA

grid point is assigned an exceedance probability function estimate which is defined at 22 different return periods ranging from 5-year to 2000-year flood depths. This makes visualization of flood exceedance estimates difficult as we can at best choose only a single value at a time to display on a map.

The ad hoc simultaneous selection of pseudo-surge values has been applied to roughly 10% of the grid points of interest in Figure 3.8. Implementation of the optimization begins with watersheds where the CPARLWR bandwidth parameter specifies a low number of local points to be used at each target point regression.

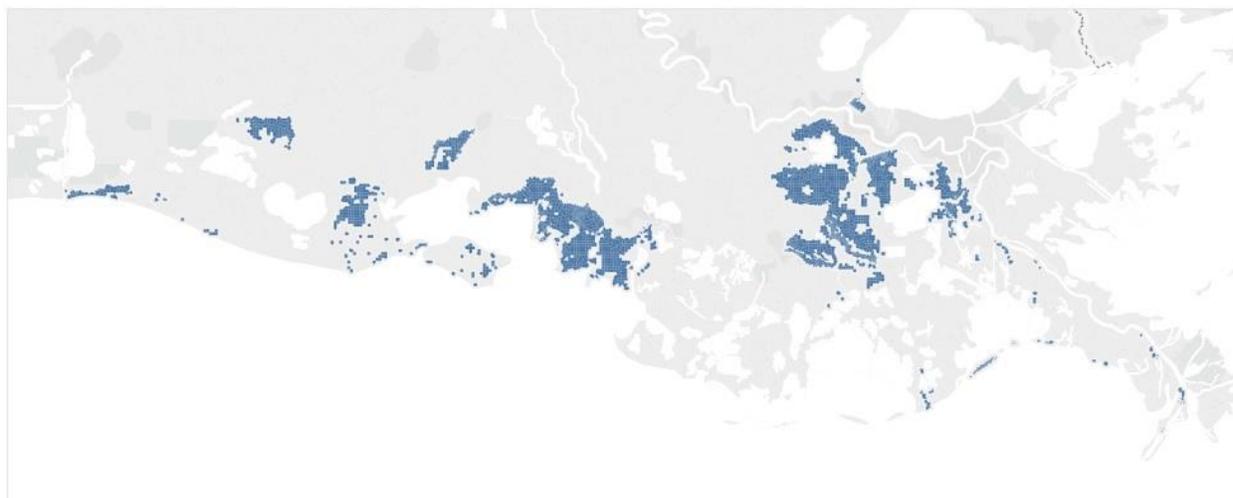


Fig. 5.1 Ad Hoc Simultaneous Selection of Pseudo-surge Data Availability

5.1 Impacts on Response Surface Performance

This section presents the impact that pseudo-surge values have on the response surface performance in terms of RMSE and storm classification rates. A storm is classified correctly if the predicted surge matches the wet or non-wet state of the simulated storm. First, results are shown for the CLARA 2017 Production, pointwise selected pseudo-surge, and topographic replacement rule pseudo-surge across the full coast.

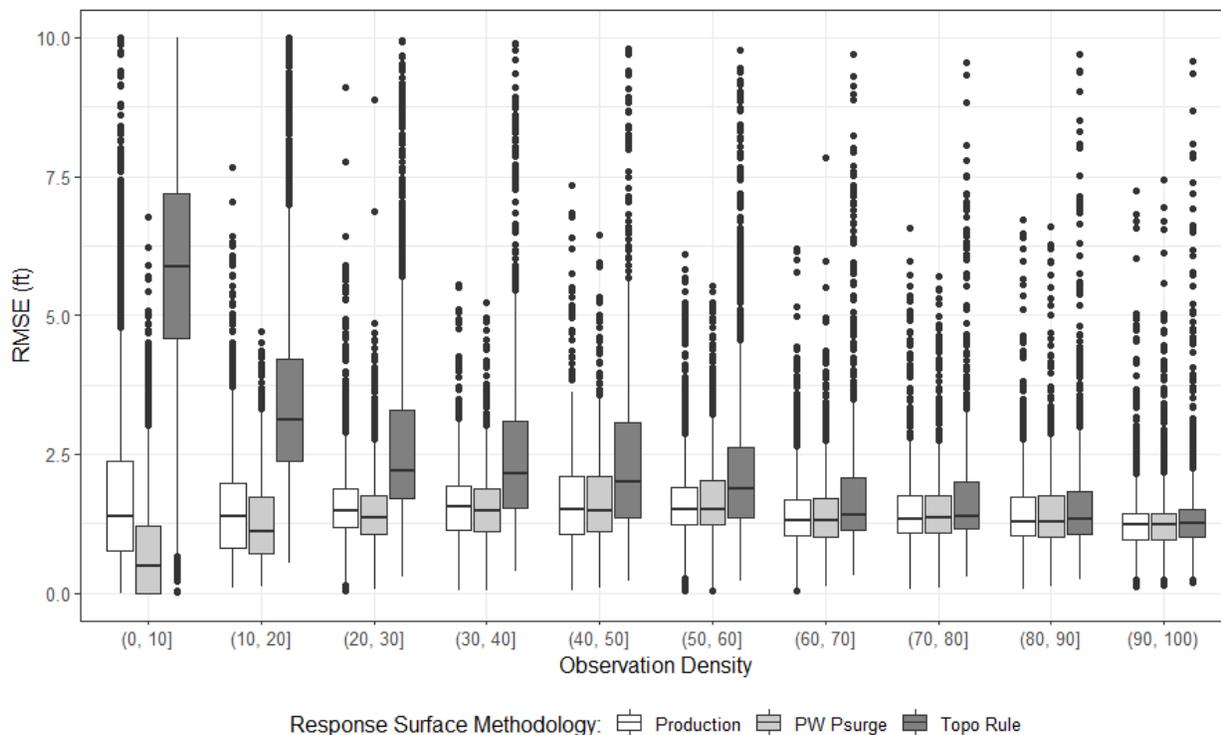


Fig. 5.2 Grid Point Wetting RMSE by Observation Density (full region)

Figure 5.2 displays the distribution of grid point RMSE aggregated across surge estimates for wetting simulation and prediction storms. The grid points are binned by their observational density. The values below each set of boxplots indicate how many grid points fall within an observational density bin. This figure compares the CLARA 2017 Production, pointwise selected pseudo-surge, and topographic replacement rule pseudo-surge response surfaces. Pointwise selected pseudo-surge values improve upon the production response surface wetting storm median RMSE for points with less than 40% observation density. Often the pointwise selected pseudo-surge implementation results in a tighter interquartile range, though an occasional

decrease in performance is observed some outlier observations. The topographic elevation replacement rule increases median bias across all grid points.

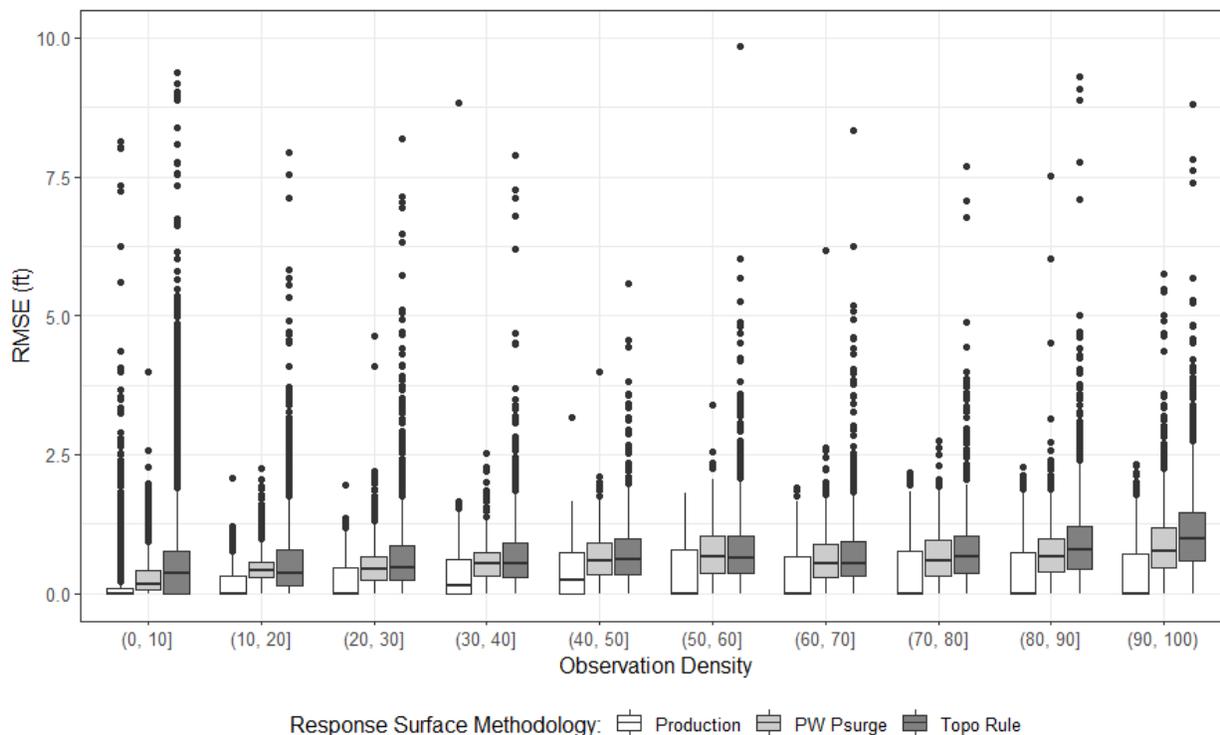


Fig. 5.3 Grid Point Non-wetting RMSE by Observation Density (full region)

Figure 5.3 displays the distribution of grid point RMSE aggregated across surge estimates for non-wetting simulation and prediction storms. Again, the grid points are binned by their observational density. For non-wetting storms, error is zero for surge estimates that are below the grid point topographic elevation. Error for all other surge estimates is measured relative to the topographic elevation. The CLARA 2017 Production response surface has a much lower median error for non-wetting storms across all observation density bins. The pointwise pseudo-surge seems to more severely misclassify non-wetting points as wetting. The classification rates of the three methods are given in Table 5.1. The CLARA 2017 Production response surface has a near perfect classification rate for non-wetting storms. Pointwise selected pseudo-surge values improve upon the RMSE for wetting storms, but do not improve upon RMSE for misclassified non-wetting storms.

Table 5.1 Storm Wet/Non-wet Classification Rates (full region)

Response Surface Method	Wet Class Rate	Non-wet Class Rate	Overall Class Rate
Production	0.938	0.990	0.981
Pointwise Pseudo-surge	0.933	0.962	0.957
Topographic Elev Rule	0.805	0.902	0.885

Next, we show results for the 10% of the coast where we have simultaneously selected pseudo-surge values via the ad hoc method in Section 4.2.4 (AHS Pseudo-surge). The region is depicted in Figure 5.1.

Table 5.2 Storm Wet/Non-wet Classification Rates (reduced region)

Response Surface Method	Wet Class Rate	Non-wet Class Rate	Overall Class Rate
Production	0.950	0.973	0.962
Pointwise Pseudo-surge	0.940	0.910	0.924
AHS Pseudo-surge	0.855	0.979	0.920
Topographic Elev Rule	0.877	0.818	0.846

The wet classification rate for the CLARA 2017 Production, pointwise pseudo-surge and topographic elevation rule are all higher for this subregion of the coast. This is likely because there is more data available. However, the non-wet classification rate drops for all methodologies. The ad hoc simultaneous pseudo-surge response surface has a low wet classification rate, but the highest non-wet classification rate. Regardless, the production response surface still performs the best overall.

Figure 5.4 shows the distribution of wetting storm RMSE aggregated at the CLARA grid point and sorted into observation density bins. For very low observation density (10-20%), the AHS pseudo-surge response surface outperforms the production version, but not as well as the pointwise pseudo-surge implementation. Grid points with between 20-30% and 50-60% display increased RMSE for the AHS method. Points with 30-50% and greater than 80% available data, the AHS pseudo-surge median grid point RMSE is high, but the interquartile range is tighter.

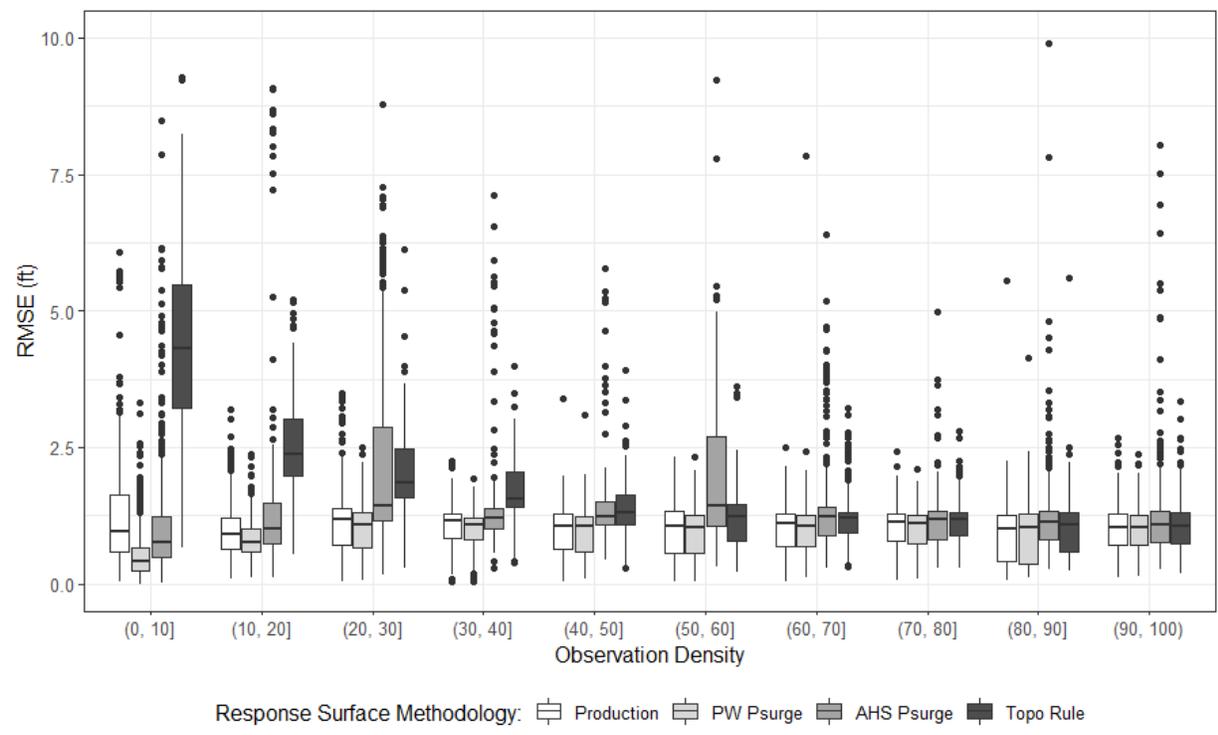


Fig. 5.4 Grid Point Wetting RMSE by Observation Density (reduced region)

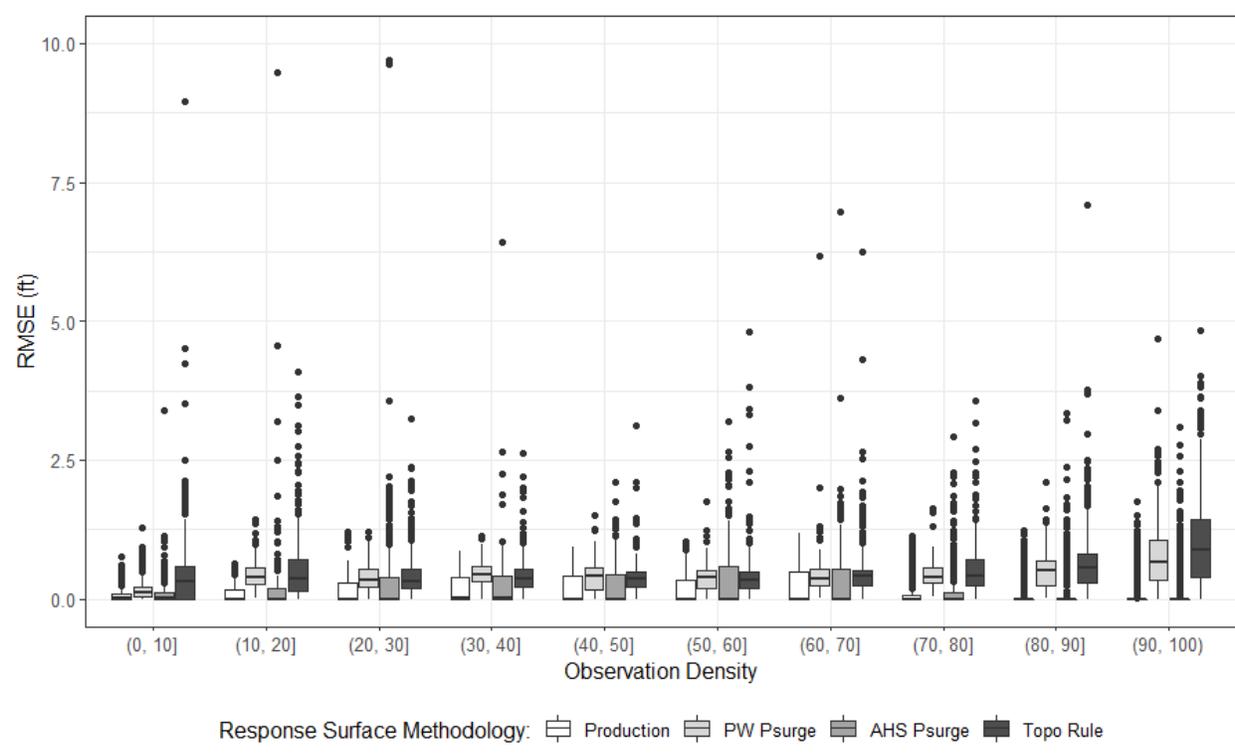


Fig. 5.5 Grid Point Non-wetting RMSE by Observation Density (reduced region)

Figure 5.5 displays the distribution of non-wet storm RMSE aggregated at CLARA grid point and sorted into observation density bins. The AHS pseudo-surge response surface performs similarly to the production version with respect to the severity of non-wetting storm misclassification. It was noted that the AHS pseudo-surge non-wet classification rate was the highest, but Figure 5.5 shows that the production response surface maintains the lowest error for misclassified non-wetting storms. Results for response surfaces with pseudo-surge values selected via the pointwise optimization or topographic elevation rule mostly follow those in Figure 5.3, though the interquartile range for pointwise pseudo-surge is much tighter for grid points with greater than 40% available data.

Recall Figure 3.9, which displays the regression method by grid point. It may be reasonable to expect the best improvement in response surface performance for grid points which were fit using a simplified regression method. Attention is given to regions where method 2 through 5 was invoked. That is (2) pointwise regression with fully specified model, (3) point-by-point regression with reduced parameters, (4) point- and track- regression, and (5) step function. With the implementation of pseudo-surge values, the best regression method these grid points can employ is the point-by-point regression with fully specified model.

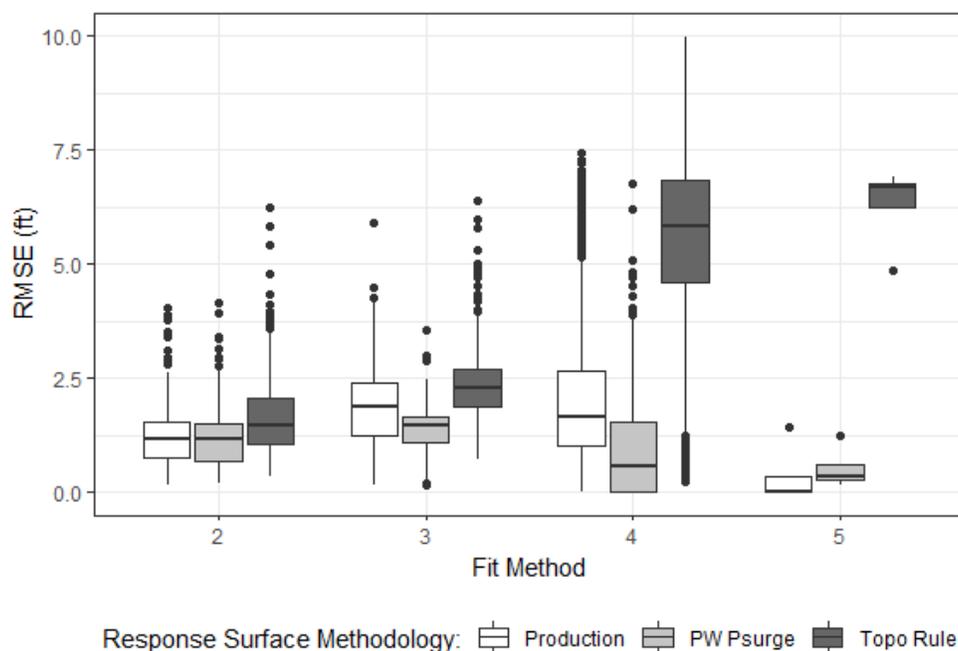


Fig. 5.6 Grid Point Wetting RMSE by Original Regression Method

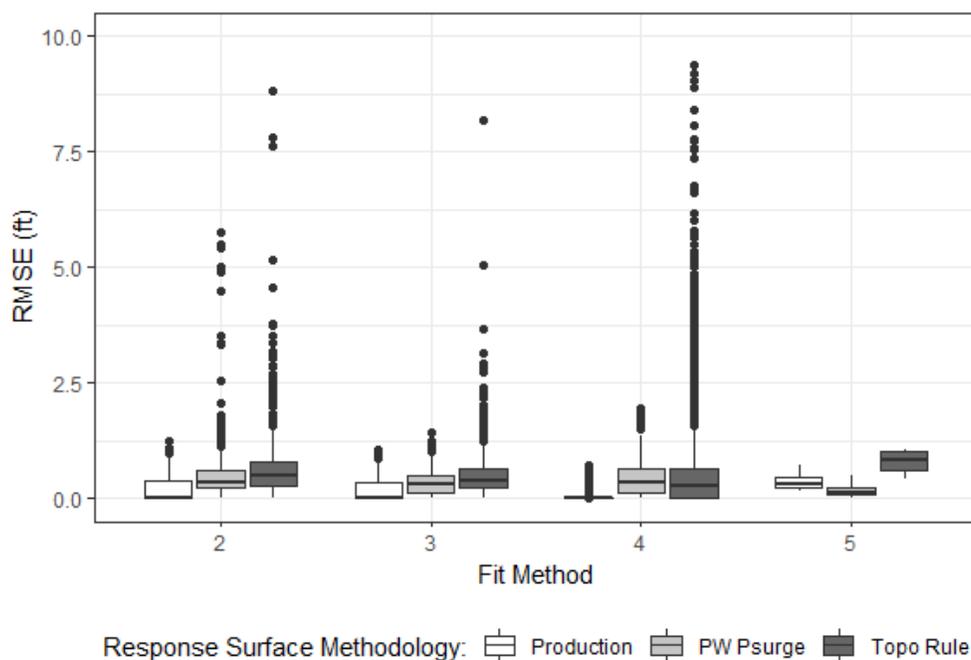


Fig. 5.7 Grid Point Non-wetting RMSE by Original Regression Method

The results in Figure 5.6 and 5.7 display the distribution of RMSE aggregated at the grid point level sorted into the four regression methods. The pointwise pseudo-surge implementation reduces error for wetting storms and increases error for non-wetting storms. It appears as if the best case of this trade-off is found for points which were fit using the reduced-form point-by-point regression. This might be expected since the reduced-form point-by-point regression simply removes some covariates from the full-form model. Note that most points are fit with the point- and track- regression method. Here, the pointwise pseudo-surge implementation does reduce error for wetting storms but does not shrink the interquartile range. Further, there is a marked increase in error for non-wetting storms.

5.2 Impacts on Flood Exceedance Curve Estimates

This section explores the impact of using pseudo-surge values on flood exceedance curve estimates. The exceedance probability function takes probabilities expressed as return periods as input and returns flood depths as output. For example, a 1000-year flood is the depth which has a 1/1000 probability of being met or exceeded in a year. The following results explore the bias

between flood depths generated with the gold standard set of 446 storms and the production, pointwise pseudo-surge, ad hoc simultaneous pseudo-surge, and topographic replacement rule implementations.

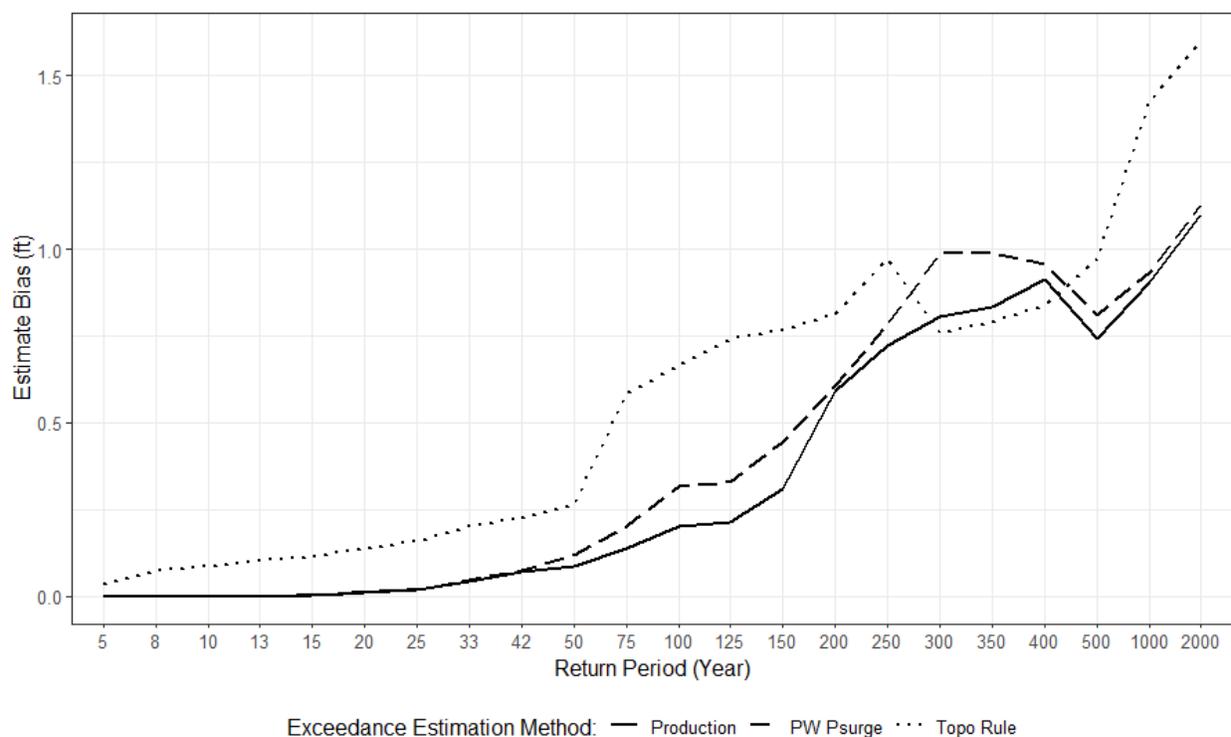


Fig. 5.8 Average Flood Exceedance Estimate Bias by Return Period (<40%) (full region)

Figure 5.8 displays the average flood depth estimate bias in feet for each of the 22 return periods across grid points with less than 40% observation density. The solid curve represents the exceedance curve estimate bias for the CLARA 2017 Production methodology while the dashed and dotted curves represent pointwise selected pseudo-surge and pseudo-surge defined with the topographic elevation substitution rule respectively. Neither the simple topographic elevation rule, nor pointwise selected pseudo-surge values reduce bias in flood exceedance estimates for any return period. However, the simple topographic elevation rule does appear to perform better in the tail of the distribution. In the previous section, the pointwise selected pseudo-surge values

reduced response surface error for wetting storms, but also increased error for non-wetting storm estimates. It appears that the result does not reduce bias in flood exceedance estimates.

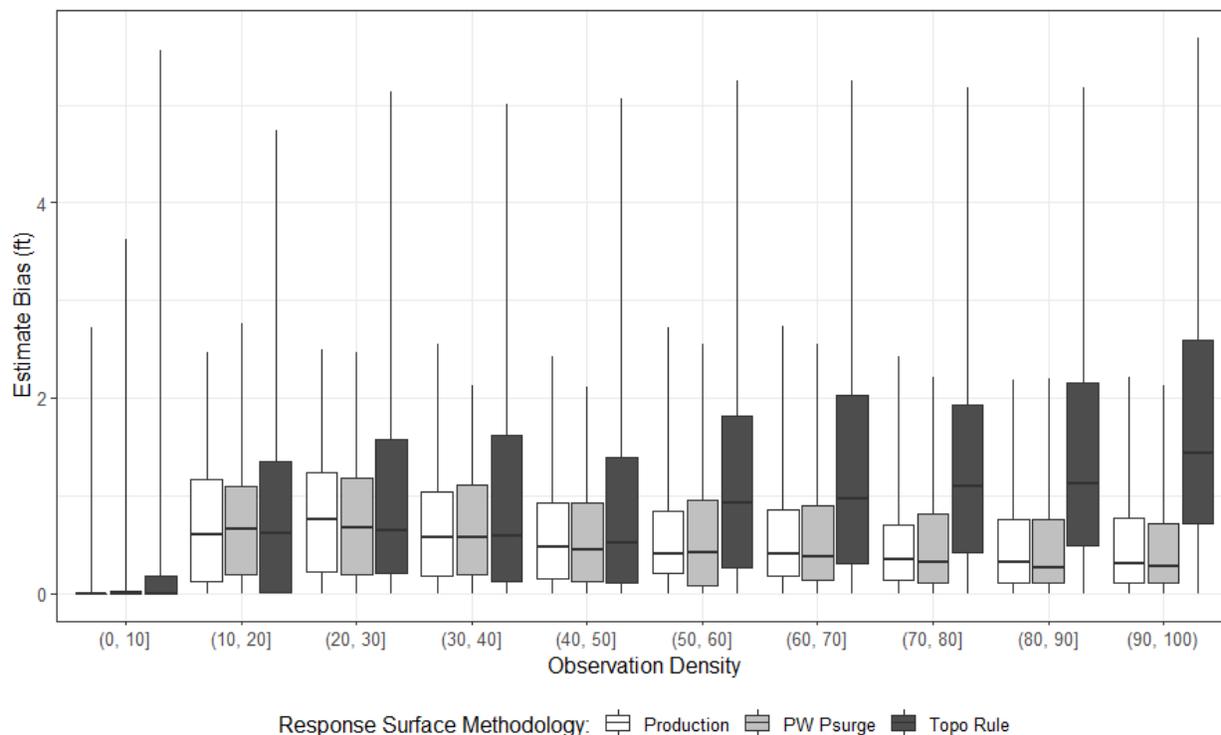


Fig. 5.9 Flood Exceedance Estimate Bias by Observation Density (100-yr) (full region)

The distribution of exceedance estimate bias aggregated at each CLARA grid point is displayed for points in each of ten observation density bins in Figure 5.9. We observe a tighter interquartile range of exceedance estimate bias when implementing pointwise selected pseudo-surge at grid points with 10-70% observation density, though the performance boost appears to be marginal at best. Again, the topographic elevation substitution rule performs unsatisfactorily across the board.

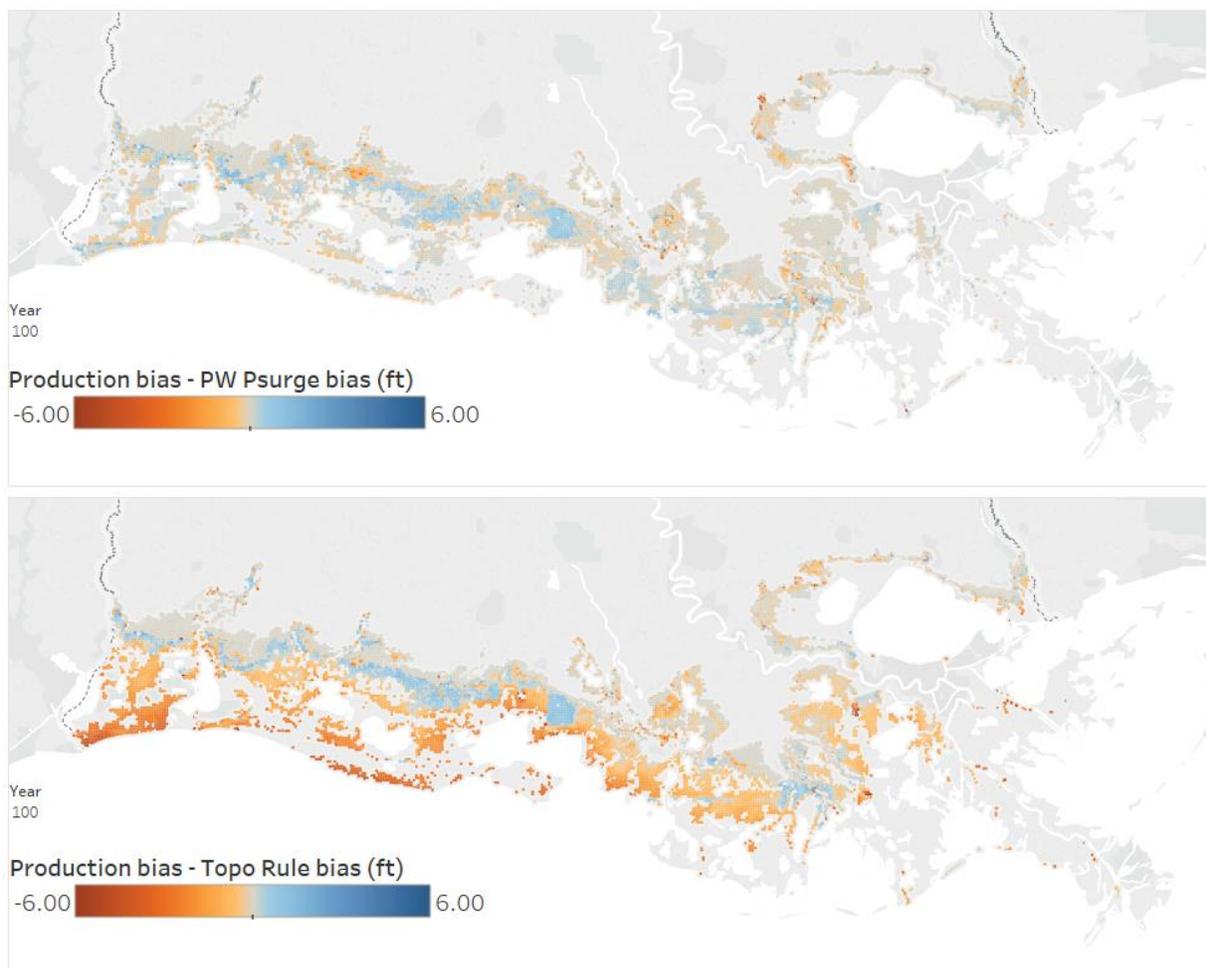


Fig. 5.10 100-year Flood Exceedance Bias Impacts (full region)

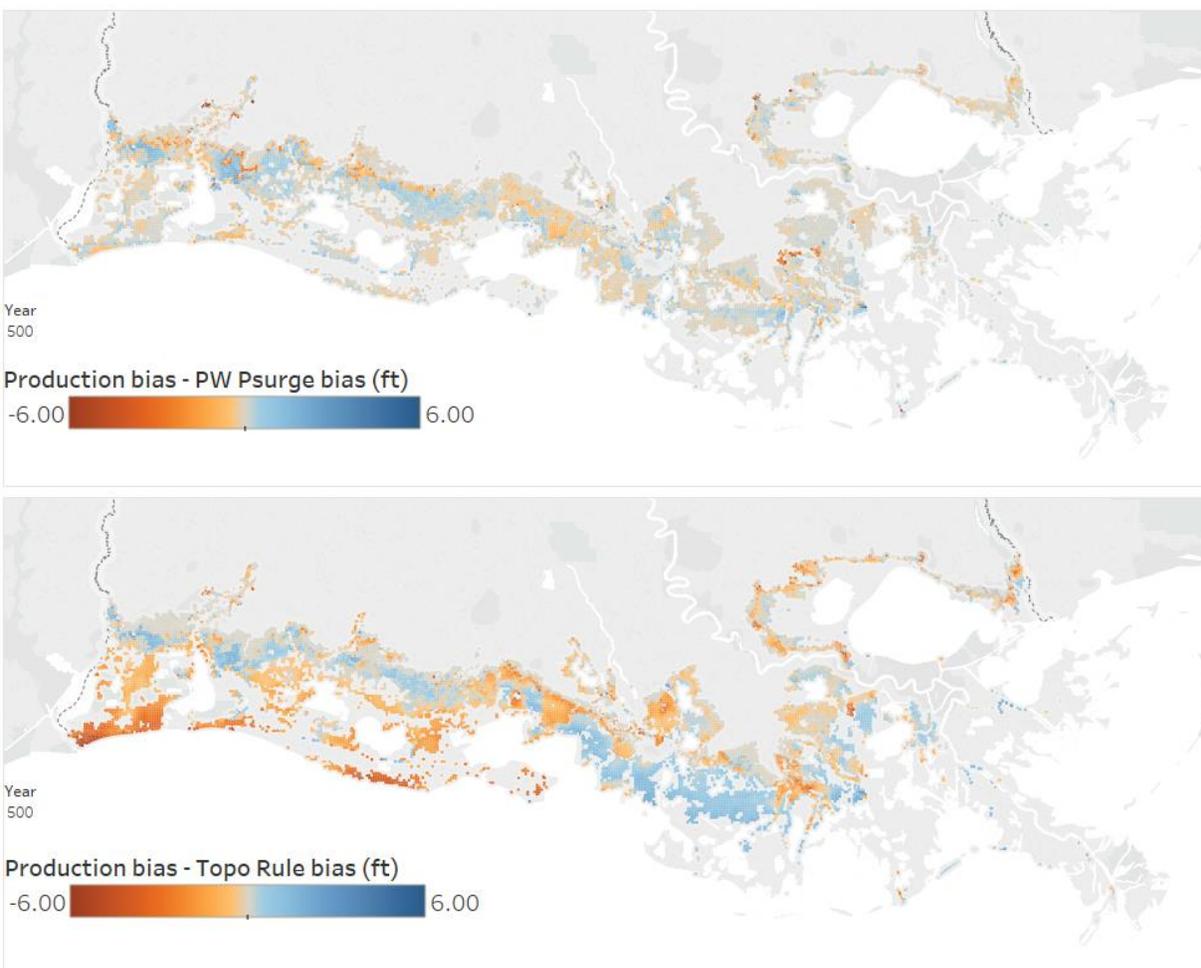


Fig. 5.11 500-year Flood Exceedance Bias Impacts (full region)

Figures 5.10 and 5.11 compare the exceedance estimates generated using pointwise selected and topographic elevation substitution rule pseudo-surge values to the 2017 CLARA production exceedance estimates. Blue and red hues indicate areas where the pseudo-surge selection method has reduced or increased, respectively, the flood exceedance estimate bias relative to the bias in the 2017 CLARA Production flood exceedance estimates. Grey areas represent little or no change in the estimate bias. A band of improvement is observed for the 100 and 500-year return periods for both pseudo-value implementations in the northwestern edge of the region. This region roughly corresponds to grid points where fewer than 40% of simulated storms caused wetting and a more simplified regression method in the response surface methodology was invoked. Though the simple topographic elevation rule negatively impacts

large tracts of other areas, it does appear to produce improved flood exceedance estimates in later return-periods. There are regions where the pointwise selected pseudo-surge values cause misclassification of some non-wetting storms in the prediction set as wetting. This results in higher estimated flood depths and overall increased bias.

Next, the current data available for the ad hoc simultaneous selection of pseudo-surge values is displayed.

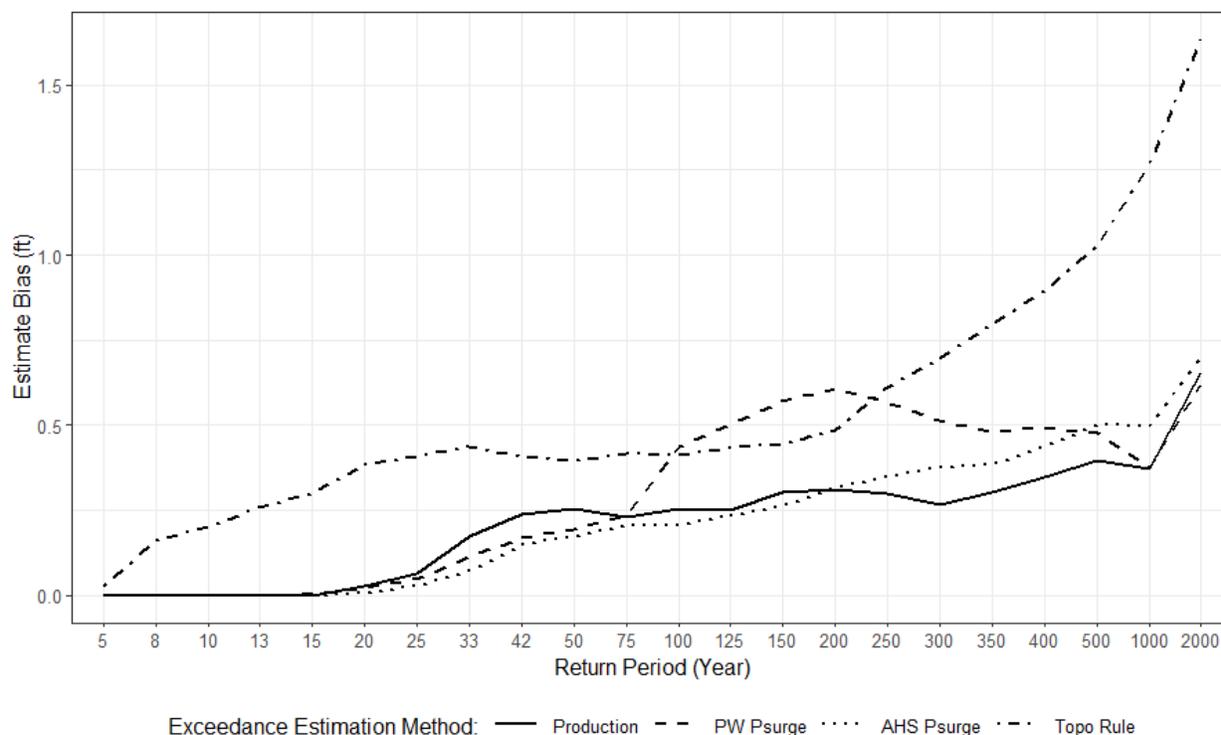


Fig. 5.12 Flood Exceedance Estimate Bias by Return Period (<50%) (reduced region)

The average flood exceedance estimate bias across all grid points for each of the twenty-two return periods is displayed in Figure 5.12. The blue curve represents the ad hoc simultaneously selected pseudo-surge. As seen here, the estimate bias is not substantially reduced for either pointwise or ad hoc simultaneously selected pseudo-surge methods. While the average bias in exceedance estimates is reduced for the 15- through 200-year return periods, performance suffers for the 200- through 2000-year return periods. Not surprisingly, the topographic elevation replacement rule performs much worse for all return periods besides the 2000-year where it only just edges out other pseudo-surge methods and the production response surface.

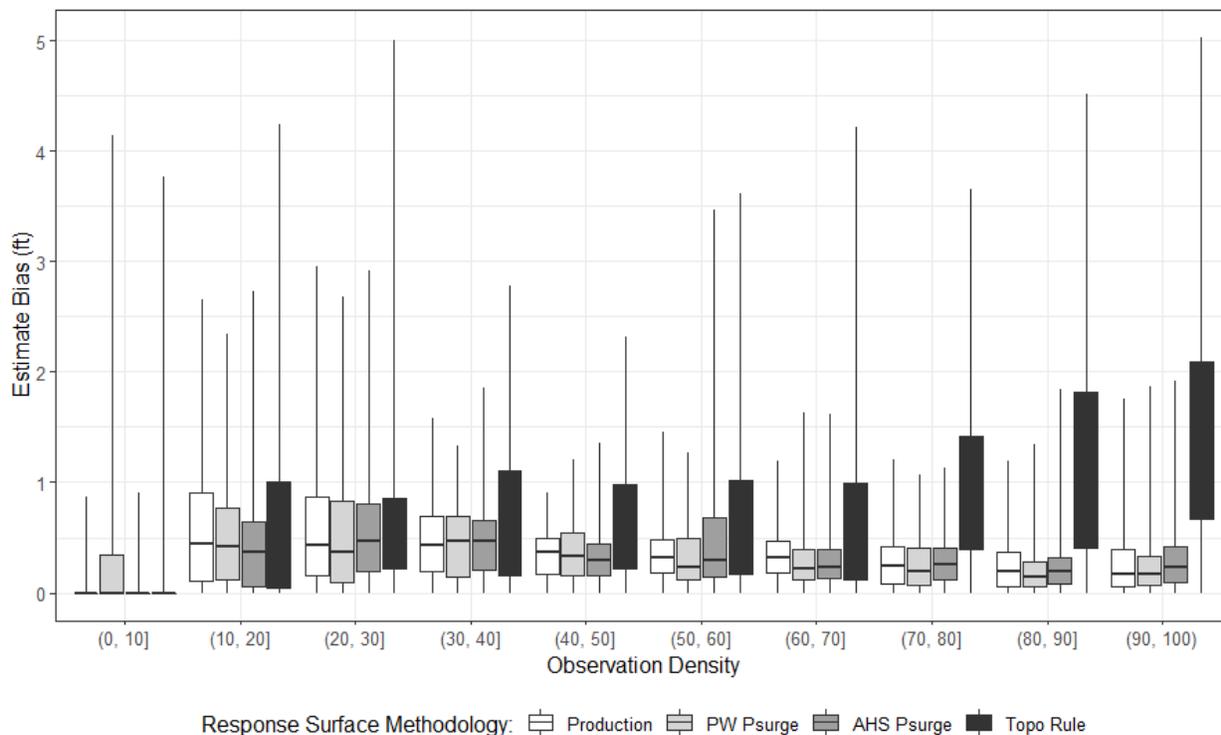


Fig. 5.13 Flood Exceedance Estimate Bias by Observation Density (100-yr) (reduced region)

Figure 5.13 displays the distribution of 100-yr flood depth estimate bias at each grid point is displayed for points in each of ten observation density bins. Here, we observe that the pointwise selected pseudo-surge generally improves median exceedance estimate bias for points with between 10-30% available observations. Slight improvement on median bias is observed for 40-80% observation density. However, pointwise selected pseudo-surge performs exceptionally poorly for points which have between 0-10% observation density. The ad hoc simultaneous method for some observation density bins though no consistent pattern is obvious. This method also tightens the interquartile range all cases except for grid points with 50-60% wetting storms. However, it is not desirable to increase the 1st quartile since this is a distribution of absolute bias for grid points. A lower bias is always preferred. Furthermore, the ad hoc simultaneous method appears to suffer from an increased number of outlier points with very high bias in 100-yr flood depths. The reasons behind this inconsistent performance are discussed in Section 5.3.

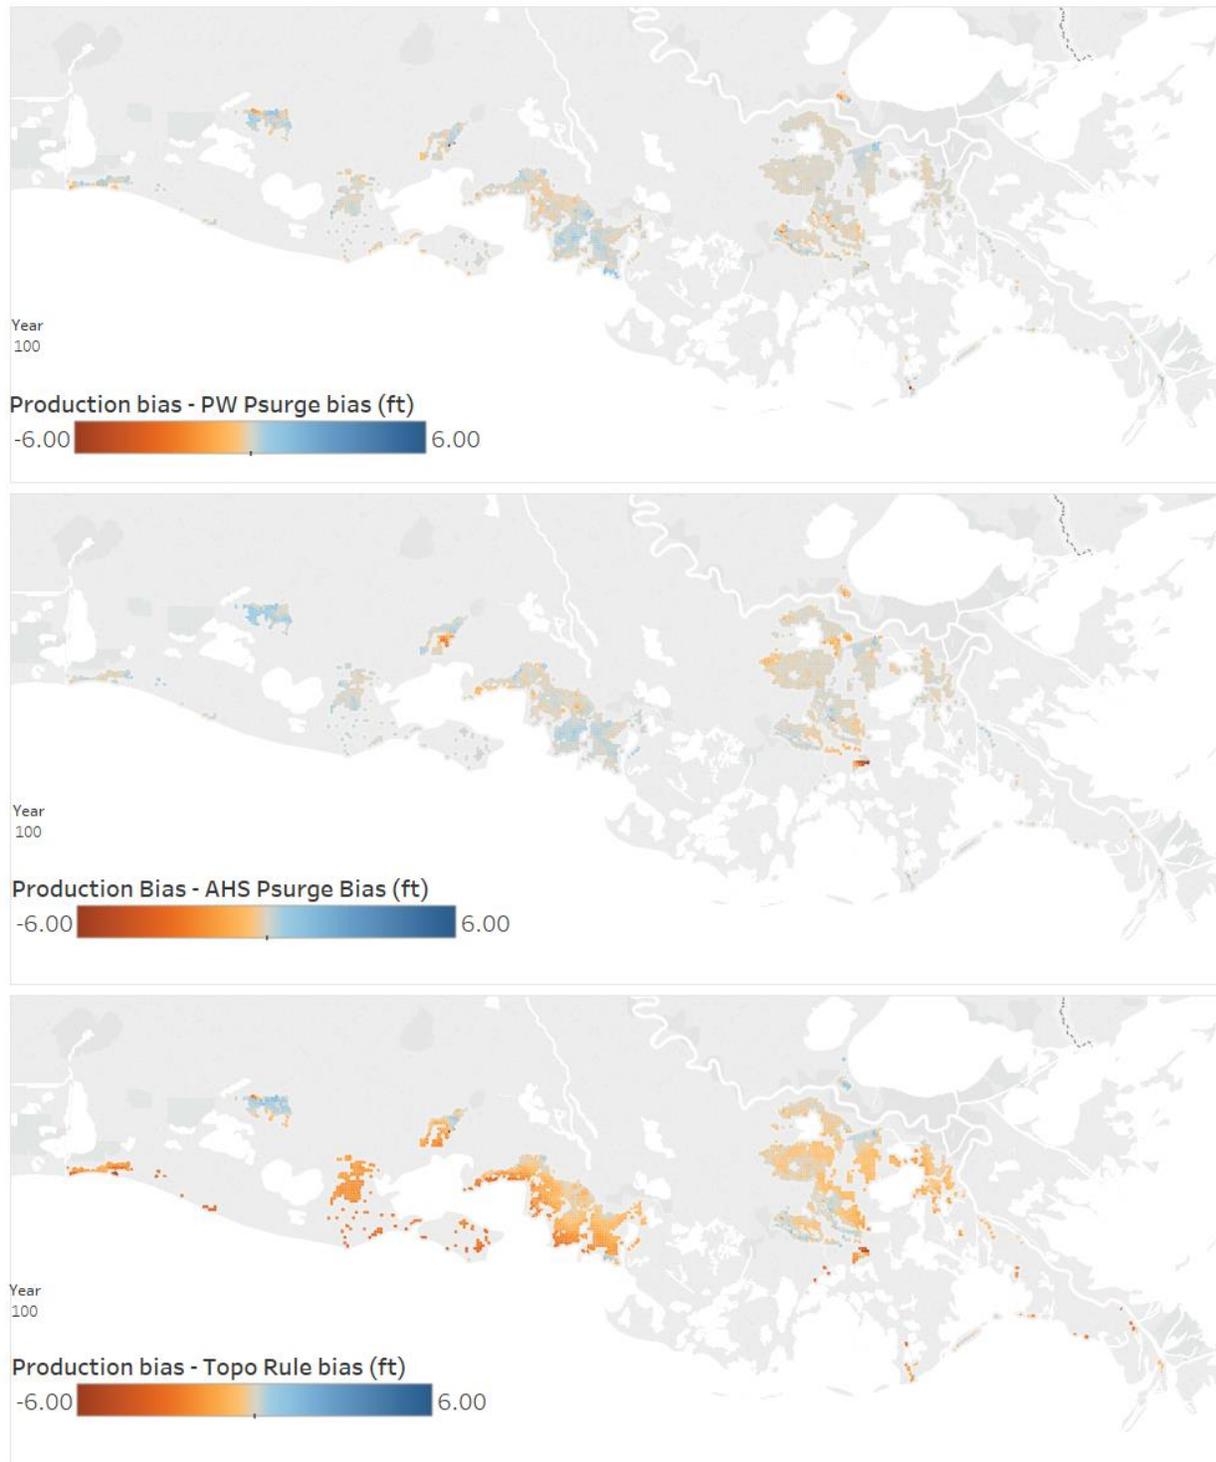


Fig. 5.14 100-year Flood Exceedance Bias Impacts (reduced region)

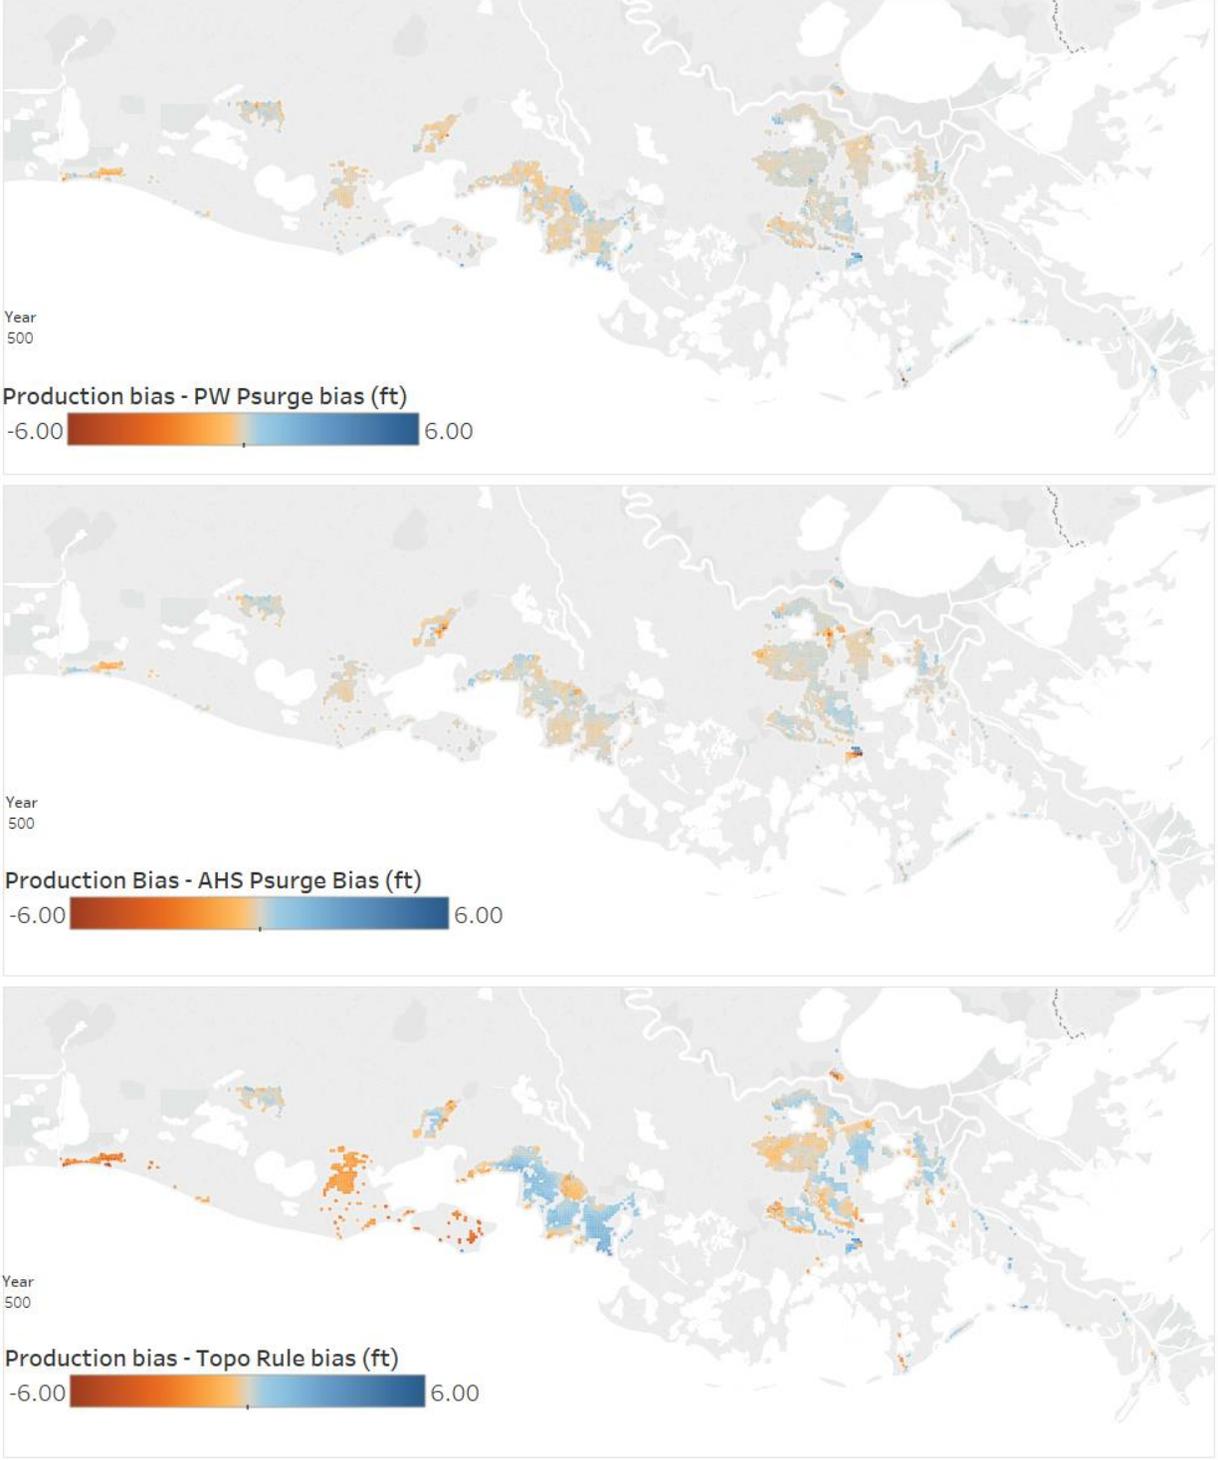


Fig. 5.15 500-year Flood Exceedance Bias Impacts (reduced region)

The maps in Figure 5.14 and Figure 5.15 compare the CLARA 2017 Production average flood exceedance bias to pointwise selected, ad hoc simultaneously selected, and topographic elevation substitution rule pseudo-surge implementations in the reduced region for the 100 and 500-year return periods. Blue and red hues indicate areas where the pseudo-surge selection method has reduced or increased, respectively, the average flood exceedance estimate bias relative to the average bias in the 2017 CLARA Production flood exceedance estimates. Grey areas represent little or no change in the average estimate bias.

As expected, the pointwise and simultaneous selection methods produced pseudo-surge values which have little to no impact across the region. At the 500-year return period the simple rule defining pseudo-surge as a function of topographic elevation reduced average estimate bias for a substantial region, but also severely increased estimate bias in most other regions resulting in an overall increase in average estimate bias as noted in Figure 5.12. At the 100-year return period the simple topographic rule increases average bias in nearly all regions. A similar effect is observed for the 2000-year return period. Though such a large region has a reduced average estimate bias, using the simple topographic elevation substitution rule may negatively impact estimates at too many other grid points for too many return periods.

5.3 Discussion

The analysis supports that pointwise pseudo-surge does appear to improve flood depth exceedance estimates for grid points where less than 40% of simulated storms caused wetting. This improvement was observed mostly on the western half of the study region. Improvements appear to coincide with regions where the response surface had previously been trained using a low-ranking regression method from the response surface hierarchy. This is also consistent with the result that the response surface was most improved for points where regression method (4) was invoked. However, no claim can be made that any pseudo-surge implementation uniformly improves flood exceedance curve estimates. Some possible explanations for this are now discussed.

It may be possible to improve the optimization problem formulation or implementation. The objective function seeks to minimize error for wetting storms in the simulation and prediction sets as well as to correctly classify non-wetting storms in the prediction set. Originally, non-wetting storms in the simulation set that are predicted to be wetting by the

response surface are reclassified to be non-wetting regardless of the prediction. However, the normal equations that produce the regression coefficients as a function of the surge and pseudo-surge values are minimizing the error non-wet storms in the simulation set as well. This is a potential source for additional regression error.

Using the normal equations to produce regression coefficients can lead to computationally near-singular matrices. Reciprocal condition numbers are tracked and matrix inverse operations are calculated outside of the non-linear optimization algorithm itself to mitigate this issue.

Though there is evidence that a linear regression model is adequate to produce storm surge estimates as a function of storm parameters (Resio, 2007; Fischbach et al., 2017), though some areas of the coast may have sufficiently complex coastal topography and bathymetry to render such a model insufficient.

The ad hoc approach to simultaneously selecting pseudo-surge values described in Section 4.2.5 may produce sub-optimal results because the method uses a weighted average of many selected pseudo-surge values. When an undefined observation is treated as local for two or more CLARA grid points, there is a possibility that the ad hoc simultaneous optimization could pick very different pseudo-values. After a weighted average of all pseudo-values is calculated, the final value may be a poor choice of pseudo-surge with respect to any target point regression.

The CLARA 2017 Production response surface methodology makes use of a post-processing heuristic to reclassify some storms in the prediction set as non-wetting even if the response surface predicts that the storms would wet. For each non-wetting storm in the simulation set, storms as or less severe in the prediction set are set to be non-wetting, regardless of the response surface prediction³. This may explain how the production response surface has a near-perfect classification rate for non-wetting storms and why optimal pseudo-surge values have little effect on flood exceedance estimates.

The error due to misclassified prediction set storms is minimized in the pseudo-surge selection optimization even though many storms will be ultimately reclassified as non-wetting by the post-processing heuristic. The optimization objective function could be reformulated in order to ignore any misclassification of storms that the post-processing heuristic will reclassify, or the

³ A storm is considered as or less severe than another storm if its central pressure is 30 mbars higher, or if its central pressure is higher, its radius of maximum windspeed is smaller, and its forward velocity is faster.

post-processing heuristic could be removed from the response surface methodology altogether and compare performance flood exceedance estimates. The pointwise pseudo-surge optimization method achieves similar average accuracy without the need for a somewhat arbitrary cleaning heuristic. The performance of the production and pointwise pseudo-surge response surfaces with and without the post-processing heuristic is examined in APPENDIX D.

It is possible that the post-processing heuristic may have affected the decision to use a 92-storm simulation set. There is no variation in forward velocity for all 168 storms used to construct flood exceedance curves. This means that the conditions for the post-processing heuristic to reclassify a misclassified storm are more easily satisfied. Any other simulation subset of the 446 storms may not have such homogenous non-wetting. The post-processing heuristic produces desired results for the 92-storm simulation set and current conditions, but its performance is unknown for other future states of the world modeled in CLARA. Though a similar statement can be made for pseudo-surge values, their use may be preferable to the current heuristic when controlling prediction set storm misclassification and its effect on flood exceedance curve estimates. One way to determine if this is the case would be to simulate all 446 synthetic storms as a reference set for a state of the world other than the current conditions.

Finally, summary statistics for response surface error and flood exceedance estimate bias that are aggregated at CLARA grid points and at different return periods. Another point of discussion is the effect on the uncertainty regarding individual storm surge estimates. Including pseudo-surge values increases the effective sample size of the regression, which leads to a decrease in estimate standard error. However, since all pseudo-surge values are required to be below the topographic elevation, response data outside the original distribution of surge is added. Therefore, sample variance will also increase. Depending on the magnitude of the pseudo-surge values, it is possible that increased sample variance outweighs the benefit of the increased effective sample size.

6. CONCLUSIONS AND FUTURE WORK

Ultimately, the goal is to improve the procedure which estimates the flood depth exceedances across the Louisiana coast. When there is a sufficiently low probability that a given location becomes inundated, the hydrodynamic simulation data becomes limited to the point where statistical models may no longer be identifiable. To avoid these cases, this thesis proposes pseudo-surge values be used for simulation data that is undefined thus increasing the effective sample size when training a response surface model. This thesis aims to investigate the potential value of information that optimal pseudo-surge values could provide.

The CLARA model is used to assess risk for many future states of the world. In these cases, a rule of thumb is required to define good pseudo-surge values for non-wetting storms in the simulation set. Such a rule might include things like 1) the topographic elevation at the location where surge is undefined, 2) the distance to the nearest wetting location(s) under the same storm conditions, 3) the elevation of the nearest wetting location(s) and, 4) the surge at the nearest wetting location(s). The methods presented here are intended to determine the best possible pseudo-surge values to provide a baseline from which a generalized rule of thumb could be vetted. However, coastal topography and climate conditions are unique, thus it is recommended that the use of a pseudo-surge rule be restricted to the geographic domain from which it was derived. Flood risk studies performed in other coastal domains could develop their own domain-specific pseudo-surge rule by simulating a large reference set of synthetic storms under current conditions from which to test response surface methodologies using various “simulation” storm subsets.

Past implementations included a simple rule which defined pseudo-surge values as a function of the topographic elevation at a point only. While this may lead to improved flood exceedance curve estimates in some locations, it produces very poor exceedance estimates in many others and is tied to an arbitrary elevation datum. It also treats all non-wetting storms as equal, which is what we wish to avoid. The attempt to find the optimal pseudo-surge values using a non-linear optimization approach resulted in a marginal reduction in exceedance estimate bias and only for some return periods.

Through the techniques discussed here, there do not appear to be pseudo-surge values which are able to uniformly improve flood exceedance curve estimates. When viewing results

geographically, we see that areas further inland and in the western half of the state may benefit from some implementation of pseudo-surge. Benefits are noted in this region at CLARA grid points for which <40% of storms in the simulation set are wetting. Though differences in flood depth exceedance estimates appear small, sometimes less than 0.1 ft, there can be comparatively large effects in estimates of expected annual damages due to storm hazards generated via the CLARA model (Johnson et al., 2018). We do not wish to harshly discount the potential value of a pseudo-surge implementation.

One interesting point of interest is that the CLARA 2017 production flood exceedance estimates are highly sensitive to a post-processing storm reclassification heuristic designed to reclassify misclassified non-wetting storms in the prediction set as non-wetting. Removing the post-processing heuristic significantly increases bias of the CLARA 2017 production exceedance estimates, but only slightly increases bias of exceedance estimates generated with the pointwise pseudo-surge implementation. It may be preferable to implement a pseudo-surge value approach in place of the post-processing heuristic if it can be shown that one method produces improved results for many future states of the world or with many different storm simulation subsets. If it is determined that pseudo-surge is preferable, then the next stage of the research is to develop a rule of thumb for including pseudo-surge in flood exceedance estimates for future states of the world.

Considering the end goal, using pseudo-surge values in the state-of-the-art response surface methodology is only one path to improve the regression scheme for hydrodynamic simulation output. There is merit for pursuing an implementation of pseudo-surge values, though a generalized rule of thumb, potentially to be used in other flood risk studies, may be ill-advised. Instead, domain-specific rules of thumb should be developed.

Coastal topography, bathymetry and weather conditions are complex. We should not expect improvements to come easily, but this thesis attempts to leverage more simulated storm information than what is currently used. Improvements to the response surface performance will produce better flood exceedance estimates, which would lead to better estimates of expected annual damage due to storm hazards, which would help stakeholders in Louisiana make more informed decisions about which risk reduction projects to implement.

APPENDIX A. JPM-OS SYNTHETIC STORM DEFINITIONS

ID	C _p	R _{max}	θ	Track	Sim
1	960	11	0	E1	Yes
2	960	21	0	E1	No
3	960	35.6	0	E1	Yes
4	930	8	0	E1	No
5	930	17.7	0	E1	Yes
6	930	25.8	0	E1	No
7	900	6	0	E1	Yes
8	900	14.9	0	E1	No
9	900	21.8	0	E1	Yes
10	960	11	0	E2	Yes
11	960	21	0	E2	No
12	960	35.6	0	E2	Yes
13	930	8	0	E2	No
14	930	17.7	0	E2	Yes
15	930	25.8	0	E2	No
16	900	6	0	E2	Yes
17	900	14.9	0	E2	No
18	900	21.8	0	E2	Yes
19	960	11	0	E3	Yes
20	960	21	0	E3	No
21	960	35.6	0	E3	Yes
22	930	8	0	E3	No
23	930	17.7	0	E3	Yes
24	930	25.8	0	E3	No
25	900	6	0	E3	Yes
26	900	14.9	0	E3	No
27	900	21.8	0	E3	Yes
28	960	11	0	E4	Yes
29	960	21	0	E4	No
30	960	35.6	0	E4	Yes
31	930	8	0	E4	No
32	930	17.7	0	E4	Yes
33	930	25.8	0	E4	No
34	900	6	0	E4	Yes
35	900	14.9	0	E4	No

ID	C _p	R _{max}	θ	Track	Sim
208	900	14.9	0	W1	No
209	900	21.8	0	W1	Yes
210	960	11	0	W2	Yes
211	960	21	0	W2	No
212	960	35.6	0	W2	Yes
213	930	8	0	W2	No
214	930	17.7	0	W2	Yes
215	930	25.8	0	W2	No
216	900	6	0	W2	Yes
217	900	14.9	0	W2	No
218	900	21.8	0	W2	Yes
219	960	11	0	W3	Yes
220	960	21	0	W3	No
221	960	35.6	0	W3	Yes
222	930	8	0	W3	No
223	930	17.7	0	W3	Yes
224	930	25.8	0	W3	No
225	900	6	0	W3	Yes
226	900	14.9	0	W3	No
227	900	21.8	0	W3	Yes
228	960	11	0	W4	Yes
229	960	21	0	W4	No
230	960	35.6	0	W4	Yes
231	930	8	0	W4	No
232	930	17.7	0	W4	Yes
233	930	25.8	0	W4	No
234	900	6	0	W4	Yes
235	900	14.9	0	W4	No
236	900	21.8	0	W4	Yes
237	960	11	0	W5	Yes
238	960	21	0	W5	No
239	960	35.6	0	W5	Yes
240	930	8	0	W5	No
241	930	17.7	0	W5	Yes
242	930	25.8	0	W5	No

ID	C _p	R _{max}	θ	Track	Sim
36	900	21.8	0	E4	Yes
37	960	11	0	E5	Yes
38	960	21	0	E5	No
39	960	35.6	0	E5	Yes
40	930	8	0	E5	No
41	930	17.7	0	E5	Yes
42	930	25.8	0	E5	No
43	900	6	0	E5	Yes
44	900	14.9	0	E5	No
45	900	21.8	0	E5	Yes
46	960	18.2	-45	E1	Yes
47	960	24.6	-45	E1	Yes
48	900	12.5	-45	E1	No
49	900	18.4	-45	E1	No
50	960	18.2	-45	E2	Yes
51	960	24.6	-45	E2	Yes
52	900	12.5	-45	E2	No
53	900	18.4	-45	E2	No
54	960	18.2	-45	E3	Yes
55	960	24.6	-45	E3	Yes
56	900	12.5	-45	E3	No
57	900	18.4	-45	E3	No
58	960	18.2	-45	E4	Yes
59	960	24.6	-45	E4	Yes
60	900	12.5	-45	E4	No
61	900	18.4	-45	E4	No
66	960	18.2	45	E1	Yes
67	960	24.6	45	E1	Yes
68	900	12.5	45	E1	No
69	900	18.4	45	E1	No
70	960	18.2	45	E2	Yes
71	960	24.6	45	E2	Yes
72	900	12.5	45	E2	No
73	900	18.4	45	E2	No
74	960	18.2	45	E3	Yes
75	960	24.6	45	E3	Yes
76	900	12.5	45	E3	No
77	900	18.4	45	E3	No

ID	C _p	R _{max}	θ	Track	Sim
243	900	6	0	W5	Yes
244	900	14.9	0	W5	No
245	900	21.8	0	W5	Yes
401	975	11	0	W1	No
402	975	21	0	W1	Yes
403	975	35.6	0	W1	No
404	975	11	0	W2	No
405	975	21	0	W2	Yes
406	975	35.6	0	W2	No
407	975	11	0	W3	No
408	975	21	0	W3	Yes
409	975	35.6	0	W3	No
410	975	11	0	W4	No
411	975	21	0	W4	Yes
412	975	35.6	0	W4	No
413	975	11	0	W5	No
414	975	21	0	W5	Yes
415	975	35.6	0	W5	No
501	975	11	0	E1	No
502	975	21	0	E1	Yes
503	975	35.6	0	E1	No
504	975	11	0	E2	No
505	975	21	0	E2	Yes
506	975	35.6	0	E2	No
507	975	11	0	E3	No
508	975	21	0	E3	Yes
509	975	35.6	0	E3	No
510	975	11	0	E4	No
511	975	21	0	E4	Yes
512	975	35.6	0	E4	No
513	975	11	0	E5	No
514	975	21	0	E5	Yes
515	975	35.6	0	E5	No
516	975	18.2	-45	E1	Yes
517	975	24.6	-45	E1	Yes
518	975	18.2	-45	E2	Yes
519	975	24.6	-45	E2	Yes
520	975	18.2	-45	E3	Yes

ID	C_p	R_{max}	θ	Track	Sim
78	960	18.2	45	E4	Yes
79	960	24.6	45	E4	Yes
80	900	12.5	45	E4	No
81	900	18.4	45	E4	No
201	960	11	0	W1	Yes
202	960	21	0	W1	No
203	960	35.6	0	W1	Yes
204	930	8	0	W1	No
205	930	17.7	0	W1	Yes
206	930	25.8	0	W1	No
207	900	6	0	W1	Yes

ID	C_p	R_{max}	θ	Track	Sim
521	975	24.6	-45	E3	Yes
522	975	18.2	-45	E4	Yes
523	975	24.6	-45	E4	Yes
524	975	18.2	45	E1	Yes
525	975	24.6	45	E1	Yes
526	975	18.2	45	E2	Yes
527	975	24.6	45	E2	Yes
528	975	18.2	45	E3	Yes
529	975	24.6	45	E3	Yes
530	975	18.2	45	E4	Yes
531	975	24.6	45	E4	Yes

APPENDIX B. CLARA WATERSHED CHARACTERISTICS

ID	Number of Points	Number of local points	Defined Obs	Undefined Obs	Total Obs	Observation Density
1	686	14	63,110	2	63,112	1.00
5	2,330	47	213,802	558	214,360	1.00
8	442	12	33,371	7,293	40,664	0.82
12	3,529	18	304,631	20,037	324,668	0.94
13	731	19	66,571	681	67,252	0.99
19	863	9	50,321	29,075	79,396	0.63
20	400	20	17,148	19,652	36,800	0.47
22	127	13	6,280	5,404	11,684	0.54
27	186	10	6,719	10,393	17,112	0.39
31	306	16	18,819	9,333	28,152	0.67
32	346	9	15,420	16,412	31,832	0.48
43	422	43	3,230	35,594	38,824	0.08
64	3,026	61	100,408	177,984	278,392	0.36
65	2,043	31	146,632	41,324	187,956	0.78
77	1,418	36	25,850	104,606	130,456	0.20
79	193	20	1,611	16,145	7,756	0.21
86	238	24	5,071	16,825	21,896	0.23
90	525	14	35,614	12,686	48,300	0.74
95	363	91	4,181	29,215	33,396	0.13
96	1,134	29	18,637	85,691	104,328	0.18
98	344	35	13,549	18,099	31,648	0.43
102	384	96	3,580	31,748	35,328	0.10
104	232	24	8,505	12,839	21,344	0.40
105	125	32	1,666	9,834	11,500	0.14
107	257	52	21,103	2,541	23,644	0.89
108	479	96	3,510	40,558	44,068	0.08
115	403	81	2,197	34,879	37,076	0.06
117	146	22	6,451	6,981	13,432	0.48
118	285	15	22,519	3,701	26,220	0.86
119	2,290	35	152,464	58,216	210,680	0.72
134	331	34	4,850	25,602	30,452	0.16
137	170	43	2,731	12,909	15,640	0.17
141	1,251	126	5,536	109,556	115,092	0.05
142	1,819	37	52,952	114,396	167,348	0.32

ID	Number of Points	Number of local points	Defined Obs	Undefined Obs	Total Obs	Observation Density
143	357	54	7,287	25,557	32,844	0.22
146	302	46	13,888	13,896	27,784	0.50
149	124	31	11,408	-	11,408	1.00
150	2,511	51	206,945	24,067	231,012	0.90
152	220	34	20,240	-	20,240	1.00
154	816	21	74,865	207	75,072	1.00
160	537	14	44,095	5,309	49,404	0.89
163	525	14	46,008	2,292	48,300	0.95
164	284	15	25,925	203	26,128	0.99
165	436	22	37,685	2,427	40,112	0.94
167	1,458	37	111,712	22,424	134,136	0.83
180	5,135	26	113,528	358,892	472,420	0.24
182	3,201	33	118,287	176,205	294,492	0.40
183	5,810	30	113,905	420,615	534,520	0.21
185	2,855	58	55,510	207,150	262,660	0.21
186	166	42	15,266	6	15,272	1.00
187	3,747	94	29,366	315,358	344,724	0.09
188	441	45	11,881	28,691	40,572	0.29

APPENDIX C. RESPONSE SURFACE METHODOLOGY DISCUSSION

This appendix describes some additional insights regarding the CLARA 2017 response surface methodology. It may be worth trying other regression methods to build the response surface. This would require testing with the full storm set and validation of the training subset approach currently used to reduce the number of simulation runs.

The current implementation of LWR defines the bandwidth parameter to capture a constant percentile of points to be considered local to a specific target point. The intent is to use a consistent number of observations to train the regression at each target point. There are two problems.

First, even though the bandwidth parameter captures an equal number of local points, that does not mean that it captures an equal number of observations. Two target points may have the same number of local points, but a very different number of defined response values. This issue could be resolved by including pseudo-surge values which “fill-in” the undefined surge response values for a given CLARA grid point. Then the LWR would use the same number of observations each time.

Second, a target point regression may be using data which is potentially thousands of meters away. This is because watersheds are often irregular shapes and sometimes sections of a region may be discontinuous. One solution could be to define the bandwidth parameter as a measure of the maximum allowable distance between the target point and local points. However, because a complete circle with radius equal to the bandwidth parameter cannot be drawn around grid points on or near a boundary, some points will lose half or more observations for training a regression.

Another issue is the variable resolution of CLARA grid points. A distance defined bandwidth parameter will define a circle that is of equal size, but a variable number of contained local points. Again, we run into the issue of unequal sample sizes. However, the reason for variable resolution in the CLARA grid points is to produce better estimates of flood effects in high population and urban areas. With a maximum distance parameter, grid points with less local information would likely be rural areas while grid points with more local information would be in urban areas. This is consistent with the purpose of the variable resolution in CLARA grid points.

The LWR method selects an optimal bandwidth parameter for each watershed via cross-validation. In the simultaneous optimization of pseudo-surge values, we use the bandwidth parameters from the CLARA2017 production response surface. In short, we are selecting pseudo-surge values which minimize the regression error with respect CLARA2017 bandwidth parameters. It is entirely possible that a different combination of pseudo-surge values and bandwidth parameters may produce more desirable results.

Fischbach et al. (2016) discussed the bias-efficiency trade-off of choosing a subset of the 446 storms to simulate. The 92-storm subset was found to have the best bias-efficiency trade-off for current conditions only.

A post-processing data cleaning heuristic appears to be largely responsible for the performance of the 2017 CLARA production response surface. This heuristic reclassifies misclassified non-wetting storms in the prediction to non-wetting by checking the surge of all simulated storms which are as or less severe. It is possible that the post-processing heuristic may have affected the decision to use a 92-storm simulation set. There is no variation in forward velocity for all 168 storms used to construct flood exceedance curves. This means that the conditions for the post-processing heuristic to reset a misclassified storm are more easily satisfied. Any other subset of the 446 storms may produce poor flood exceedance estimates if just one wetting storm as or less severe is simulated. The post-processing heuristic produces desired results for the 92-storm simulation set and current conditions, but its performance is unknown for other future states of the world modeled in CLARA.

APPENDIX D. ADDITIONAL FLOOD EXCEEDANCE ANALYSIS

Here we present additional results for the response surface performance and flood exceedance estimates generated without the post-processing heuristic discussed in Section 5.4.

Table D.1 Storm Wet/Non-wet Classification Rates (removed heuristic)

Response Surface Method	Wet Class Rate	Non-wet Class Rate	Overall Class Rate
Production	0.938	0.990	0.981
Production (RH)	0.938	0.925	0.927
Pointwise Pseudo-surge	0.933	0.962	0.957
Pointwise Pseudo-Surge (RH)	0.933	0.940	0.939

As expected, classification rates of wetting storms remain identical with and without the post-processing heuristic. However, classification rates for non-wetting storms are affected. The production response surface non-wet class rate is 0.925, while the pointwise pseudo-surge non-wet class rate is 0.940.

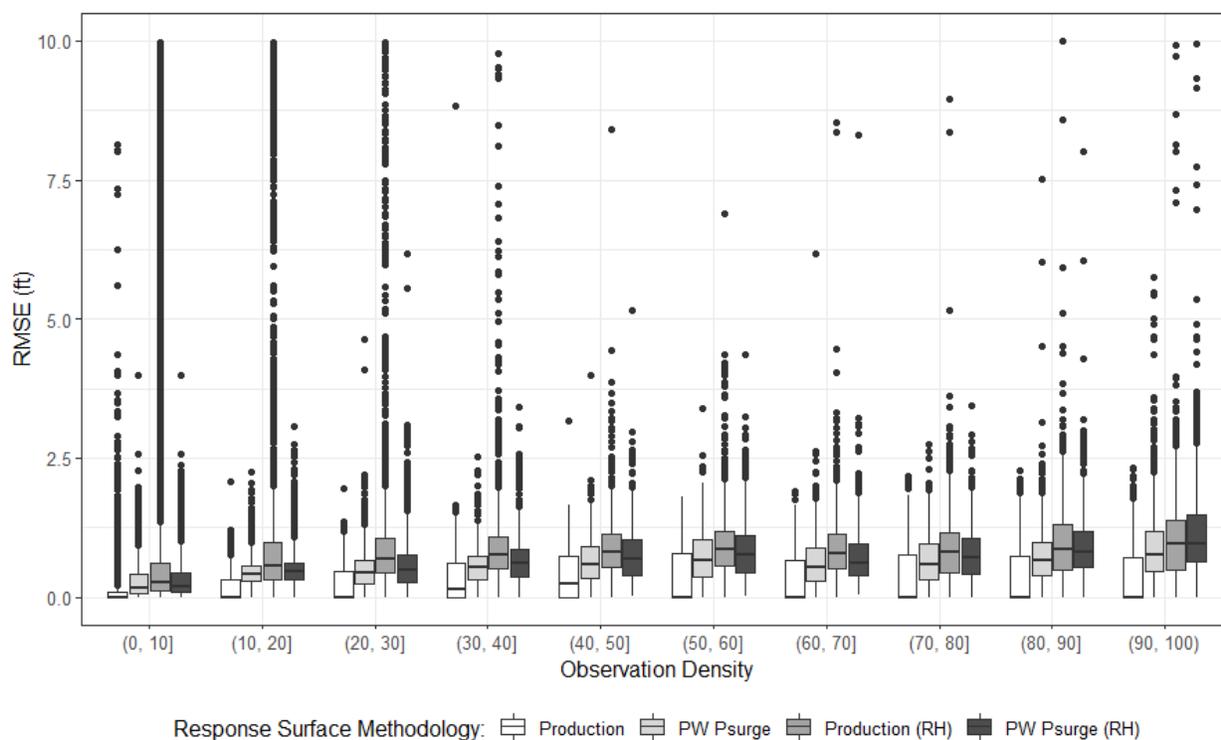


Fig. D.1 Grid Point Non-wetting RMSE by Observation Density (removed heuristic)

Figure D.1 shows the distribution of grid point RMSE for non-wetting storms. Here we observe that, with the post-processing heuristic removed, classification of non-wetting storms suffers. However, the pointwise pseudo-surge response surface maintains a much lower median RMSE for nearly all observation density bins. The pointwise pseudo-surge response surface improves upon the non-wet classification rate of the production response surface by 0.015, but also significantly reduces error for misclassified storms. The distribution of grid point wetting RMSE for is identical for response surfaces with and without the post-processing heuristic.

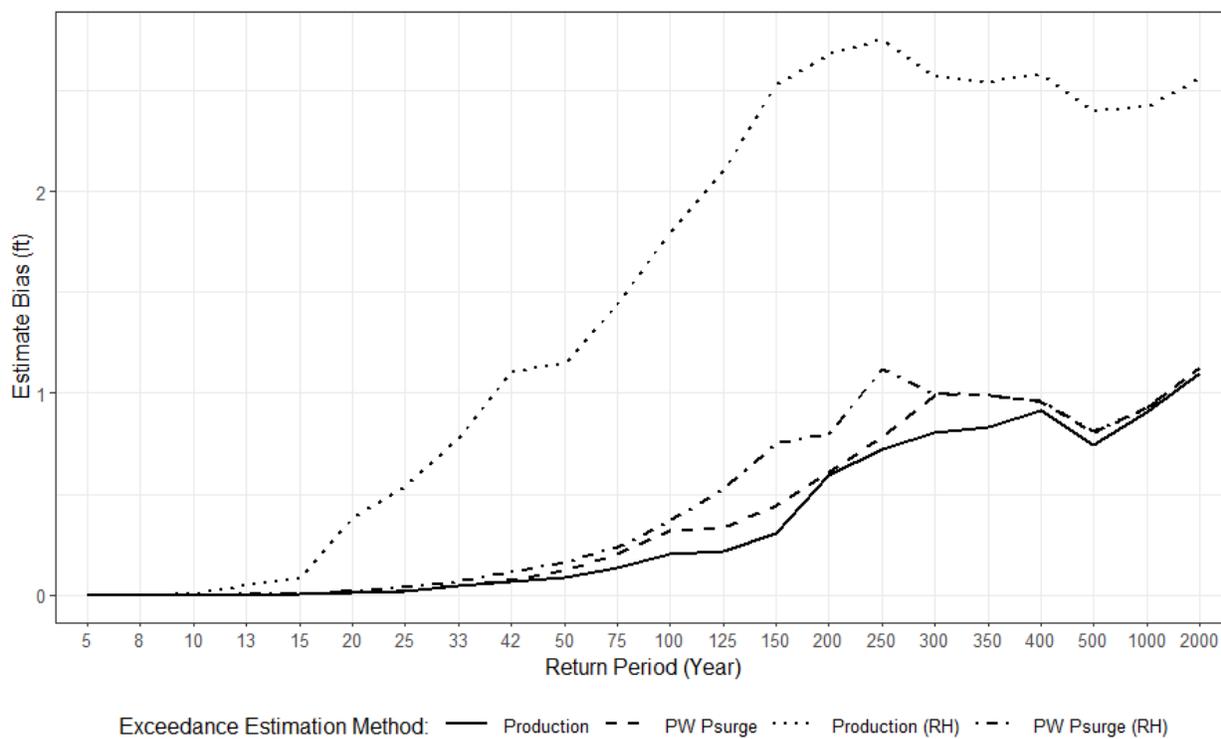


Fig. D.2 Average Flood Exceedance Estimate Bias by Return Period (removed heuristic)

Figure D.2 displays the CLARA 2017 Production and pointwise pseudo-surge response surfaces both with and without the post-processing heuristic (designated RH for ‘removed heuristic’). Without the post-processing heuristic, the 2017 CLARA Production response surface suffers from inflated flood exceedance bias estimates for nearly all return periods.

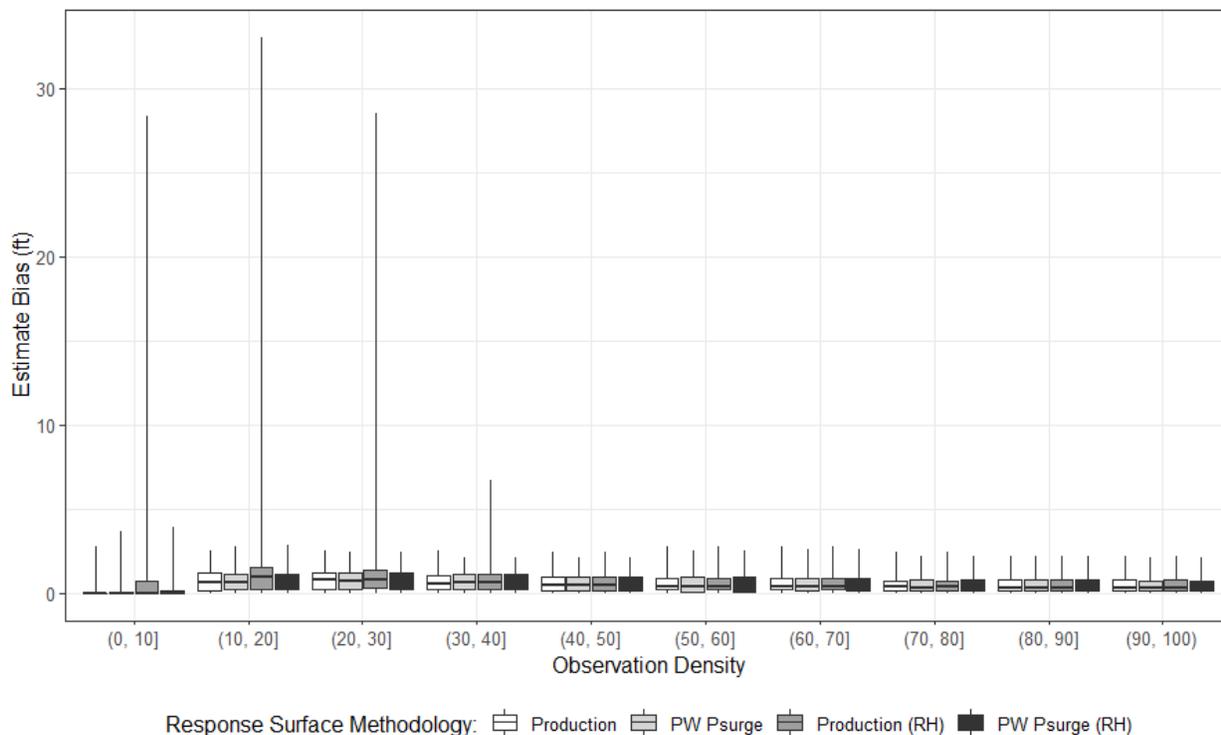


Fig. D.3 Flood Exceedance Estimate Bias by Observation Density (100-yr) (rem heuristic)

Figure D.3 shows that the CLARA 2017 Production response surface methodology without the post-processing heuristic generates poor 100-yr flood depth exceedance estimates for points with below a 30% observation density. The estimate bias is over 30 ft. with the worst performance observed in the 10-20% range. This is likely because at 10-20% observation density the response surface is just able to identify a linear regression model, but the number of training observations is so low that predicted surge values are unreliable. Once pointwise selected pseudo-surge is included, the flood exceedance estimate bias for points below 30% observation density is comparative to the bias for points above 30% observation density. This is evidence of the efficacy of pseudo-surge when the post-processing heuristic is not used.

REFERENCES

- Agbley, S. K. & Basco, D. R. (2008). An Evaluation of Storm Surge Frequency-of-Occurrence Estimators. *Solutions to Coastal Disasters Congress*.
- Dietrich, J. C., Tanaka, S., Westerink, J. J., Dawson, C. N. Luetlich, R. A., Jr. Zijlema, M., . . . Westerink, H.J. (2012). Performance of the unstructured-mesh, SWAN+ADCIRC model in computing hurricane waves and surge. *Journal of Computational Science*, 52, 468–497.
- Divoky, D. & Resio, D. T. (2007). Performance of the JPM and EST Methods in Storm Surge Studies. *10th International Workshop on Wave Hindcasting and Forecasting and Coastal Hazard Symposium*.
- Coastal Restoration and Protection Authority of Louisiana. (2012). *Louisiana's Comprehensive Master Plan for a Sustainable Coast*. Baton Rouge, LA: Coastal Restoration and Protection Authority.
- Cleveland, W. S., & Devlin, S. J. (1988). Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting. *Journal of the American Statistical Association*, 83(403), 596–610.
- Hurricane Research Division. (2014). HURDAT (Hurricane Database) [Data file]. Retrieved from http://www.aoml.noaa.cob/hrd/hurdat/Data_Storm.html.
- Luetlich, R. A., & Westerink, J.J. (2004). Formulation and Numerical Implementation of the 2d/3d Adcirc Finite Element Model Version 44.Xx. R. Luetlich: Chapel Hill, NC.
- Feuillet, T., Charreire, H., Menai, M., Salze, P., Simon, C., Dugas, J., . . . Oppert, J. (2015). Spatial heterogeneity of the relationships between environmental characteristics and active commuting: towards a locally varying social ecological model. *International Journal of Health Geographics*, 14(12), 1–14.
- Fischbach, J. R., Johnson, D. R., & Kuhn, K. (2016). Bias and Efficiency Tradeoffs in the Selection of Storm Suites Used to Estimate Flood Risk. *Marine Science and Engineering*. 4 (10).
- Fischbach, J. R., Johnson, D. R., Ortiz, D. S., Bryant, B. P., Hoover, M., & Ostwald, J. (2012). Louisiana's Comprehensive Master Plan for a Sustainable Coast, Appendix D25 – Risk Assessment (CLARA) Model Technical Report. Baton Rouge, LA. LA Coastal Protection and Restoration Authority of Louisiana.

- Fischbach, J. R., Johnson, D. R., Kuhn, K., Pollard, M., Stelzner, C., Costello, R., ... Cobell, Z. (2017). 2017 Coastal Master Plan Modeling: Attachment C3-25: Storm Surge and Risk Assessment. Version Final. (pp. 1-219). Baton Rouge, LA: Coastal Protection and Restoration Authority.
- Ho, F. P., & Myers, V. A. (1975). Joint Probability Method of Tide Frequency Analysis Applied to Apalachicola Bay and St. George Sound, Florida. NOAA Tech. Rep. NWS 18, 43pp.
- Irish, J. L., Resio, D. T., & Cialone, M. A. (2009). A surge response function approach to coastal hazard assessment Part 2: Quantification of spatial attributes of response functions. *Natural Hazards*, (51), 183–205.
- Johnson, D. R., Fischbach, J. R., & Ortiz, D. S. (2013). Estimating surge-based flood risk with the coastal Louisiana risk assessment model. *Journal of Coastal Research*, 29, 109-126.
- Johnson, D. R. (2018). Improving Methods for Estimating Flood Depth Exceedances Within Storm Surge Protection Systems. *Risk Analysis*.
- Kaplan, D., & Garrick, B. J. (1981). On the Quantitative Definition of Risk. *Risk Analysis*. 1(1).
- Koul, H., Susarla, V., & Van Ryzin, J. (1981). Regression analysis with randomly right-censored data. *The Annals of Statistics*, 9(6), 1276–1288.
- Krantz, S. G. (2015). *Convex Analysis*. New York, NY: CRC Press.
- Miller, R. G. (1976). Least squares regression with censored data. *Biometrika*, 63(3), 449–64.
- Miller, R. G., & Halpern, J. (1982). Regression with censored data. *Biometrika*, 69(3), 521–31.
- Myers, V. A. (1975). Storm Tide Frequencies on the South Carolina Coast. NOAA Tech. Rep. NWS 38, completed under agreement EMW-84-E-1589 for FEMA, 194pp.
- Nordec, J., & Wright S. J. (2006). *Numerical Optimization* (2nd ed.). New York, NY: Springer.
- Resio, D. T. (2007). White paper on estimating hurricane inundation probabilities. US Army Corps of Engineers, ERDC-CHL.
- Resio, D. T., Irish, J., & Cialone, M. (2009). A surge response function approach to coastal hazard assessment – part 1: basic concepts. In *Natural Hazards*, 51, 163-182.
- Roberts, H., & Cobell, Z. (2017). 2017 Coastal Master Plan: Attachment C3-25.1: Storm Surge. Version Final. (pp.1-110). Baton Rouge, LA: Coastal Protection and Restoration Authority.
- Rosasco, L., De Vito, E., Caponnetto, A., Piana, M., & Verri, A. (2003). Are Loss Functions All the Same? *Neural Computation*. 16(5), 1063-1076.

- Scheffner, W. N., Clausner J. E., Militello, A., Borgman, L. E., Edge, B. L. & Grace, P. J. (1999). Use and Application of the Empirical Simulation Technique: User's Guide. Technical Report CHL-99-21. U.S. Army Corps of Engineers.
- Stevens, S. S. (1946). On the Theory of Scales of Measurement. *Science*. 103(2684), 667-680.
- Thompson, M. Lou, & Nelson, K. P. (2003). Linear regression with Type I interval- and left censored response data. *Environmental and Ecological Statistics*, 10, 221–230.
- Toro, G. R., Niedoroda, A. W., Reed, C. W., & Divoky, D. (2010a). Quadrature-based approach for the efficient evaluation of surge hazard. *Ocean Engineering*, 37(1), 114–124.
- Toro, G. R., Resio, D. T., Divoky, D., Niedoroda, A. W., & Reed, C. (2010b). Efficient joint probability methods for hurricane surge frequency analysis. In *Ocean Engineering*, 37, 125-134.
- USACE (2008a). Louisiana Coastal Protection and Restoration Technical Report. Vicksburg, Mississippi: USACE, 98 pp.
- USACE (2008b). Flood Insurance Study: Southeastern Parishes, Louisiana. Intermediate Submission 2. Vicksburg, Mississippi: USACE, 697 pp.
- USACE (2009). Louisiana Coastal Protection and Restoration Technical Report: Hydraulics and Hydrology Appendix. New Orleans, LA: U.S. Army Corps of Engineers.
- Vickery, P. J., Skerlj, P. F., & Twisdale, L. A. (2000) Simulation of Hurricane Risk in the U.S. Using Empirical Track Model. In *Journal of Structural Engineering*.