Crowd-sourced Data, Data-driven Techniques, and Model Predictive Control to Improve Understanding, Prediction and Mitigation of Urban Coastal Flooding

---

A Dissertation

Presented to

the Faculty of the School of Engineering and Applied Science

University of Virginia

In Partial Fulfillment

of the requirements for the Degree

Doctor of Philosophy (Civil and Environmental Engineering)

by

Jeffrey Michael Sadler

May 2019

Abstract

Warmer average global temperatures have already caused and will likely cause further increased flooding in coastal cities. This increased flooding is due to an increase in average storm intensity and a rise in sea level, results of warmer global temperatures. In support of a coastal citys ability to adapt in place to increased flood risk, this dissertation focuses on understanding, predicting, and mitigating coastal flooding through four studies. Each of the four studies will focus on one of these three overarching objectives: to understand, to predict, and to mitigate coastal flooding. In the four studies, emerging technologies, approaches, and datasets are investigated as alternatives to more traditional approaches whose limits are pushed by complex urban coastal environments. Geographically, this dissertation focuses on two coastal cities in Southeastern Virginia USA: Norfolk and Virginia Beach, an area particularly vulnerable to coastal flooding due to land subsidence. The first study focuses on the importance of rainfall measurements close to flood-prone intersections. The objective of this study was to quantify the impact of nearby rain gauges in estimating areal rainfall for watersheds which drain to flood-prone street intersections in Virginia Beach. The second study focuses on data-driven prediction of coastal street flooding. This study demonstrates data-driven models using crowd-sourced data as a step toward street-level flood prediction. The third and fourth studies focus on the ability of active stormwater controls, particularly model predictive control (MPC) for mitigating flooding. The third study describes a software tool developed to simulate MPC of a stormwater system using the Environmental Protection Agency Stormwater Management Model (EPA-SWMM5). The fourth and final study provides insight as to how the utility of active stormwater control will change over time given different sea level rise projections and control configurations. This research found that (i) the data from a rain gauge within 0.5 km of a watershed centroid in Virginia Beach can reduce estimation variance by 50-100%, (ii) using crowd-sourced flood reports as training data, data-driven models can predict within one flood report on average for rainfall event in Norfolk, and (iii) that in addition to a tide gate, MPC can further reduce overall flooding with an average effective percent reduction of 32% in Norfolk. Norfolk and Virginia Beach are experiencing the effects of sea level rise earlier than many coastal cities. The methods developed and insight gained in this dissertation could be applied to and benefit other coastal cities who, in the coming decades, will likely experience sea level rise effects similar to those currently being experienced in Norfolk and Virginia Beach.

Approval Sheet

This dissertation is submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Civil and Environmental Engineering)

Jeffrey Michael Sadler

This dissertation has been read and approved by the Examining Committee:

Jonathan Goodall, Adviser

Teresa Culver , Committee Chair

Lawerence Band

John Porter

Madhur Behl

Accepted for the School of Engineering and Applied Science:

Craig Benson, Dean, School of Engineering and Applied Science

May 2019

To Camilla, Andrew, and Caleb

Acknowledgments

The sucessful completion of my dissertation is thanks to many who have supported me over these four years and throughout my life.

I thank my friends and co-students: Faria Zahura, Gina O'Neil, Iman Maghami, Alex Chen, Ben Bowes, Youngdon Choi, Yawen Shen, Stuart Sheffield, Ryan Mahon, Bakinam Essawy, and Linnea Saby. One of the most enriching parts of this journey has been to work, talk, walk, learn, and lunch together with you all. I especially thank Mohamed Morsy who has always made himself available to help or provide feedback since the start of my time at UVa.

I thank my PhD dissertation committee, Professors Band, Behl, Culver, Goodall, and Porter. I especially thank Jon Goodall who has provided needed guidance, enthusiasm, and encouragement throughout my PhD experience and has inspired me to be a similar mentor someday to another.

I thank my parents, who taught me how to work, who have always stressed the importance of education, and who have shown me what the important things in life are. Their confidence in me and love has made all the difference in my life.

I thank my brothers, Greg and Cory, and sisters, Kate, Patti, and Jackie for their examples and constant friendship.

I thank my boys, Andrew and Caleb. You have kept my life (and my graduate studies) in perspective. The greatest thing I will have done in graduate school is to have become your father.

I thank my wife, Camilla. You have been my constant companion and support. This degree is at least 50% yours. I will always remember with gratitude this time of our lives in which you stood by me till the end.

Finally, I acknowledge and thank my God and Savior, Jesus Christ, who has given me strength to carry my burdens and whose grace has allowed me to succeed, despite my weakness.

14 And I will also ease the burdens which are put upon your shoulders, that even you cannot feel them upon your backs, even while you are in bondage; and this will I do that ye may stand as witnesses for me hereafter, and that ye may know of a surety that I, the Lord God, do visit my people in their afflictions.

15 And now it came to pass that the burdens which were laid upon Alma and his brethren were made light; yea, the Lord did strengthen them that they could bear up their burdens with ease, and they did submit cheerfully and with patience to all the will of the Lord.

Mosiah 18:14-15, Book of Mormon

27 And if men come unto me I will show unto them their weakness. I give unto men weakness that they may be humble; and my grace is sufficient for all men that humble themselves before me; for if they humble themselves before me, and have faith in me, then will I make weak things become strong unto them.

Ether 12:27, Book of Mormon

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Warmer average global temperatures have already caused and will likely cause further increased flooding in coastal cities (Sweet and Park, 2014). This increased flooding is due to an increase in average storm intensity and a rise in sea level, results of warmer global temperatures. While projections of sea level rise vary, the Intergovernmental Panel on Climate Change predicts that sea level in 2081-2100 will have risen 0.26 - 0.98 m from 1971-2010 conditions (Church et al., 2013). Not only will higher seas increase the risk of tidal inundation, but they will decrease the capacity for coastal cities' stormwater infrastructure to drain runoff from rainfall events. This is a serious concern nationally, as 39% of the U.S. population lives near the coast (NOAA, 2017f), and internationally, as many of the global economic and population centers lie on or near the coast.

Coastal cities have several alternatives and combination of alternatives for investing resources to address this problem (Hallegatte, 2009). Perhaps the most extreme and definite solution is abandoning coastal cities altogether and retreating inland to higher elevations. This may, depending on the actual magnitude of sea level rise, be the only viable alternative in the long-term. Another alternative is to invest in built infrastructure such as flood walls and sea gates to protect against flooding. A third approach for addressing increased coastal flooding is adapting in place. In contrast to building new and bigger infrastructure, adapting in place involves retrofitting and actively managing existing infrastructure fully utilize its capacity. This would involve investing in better understanding, more accurate prediction, and the ability to proactively mitigate the impacts of flooding events.

Adapting in place to coastal flooding, compared to the abandonment and building gray infrastructure, is less straightforward and requires more innovation, however, it would likely incur less monetary and social costs, at least in the short-term (Cooper and Lemckert, 2012; Rosenzweig et al., 2011). There are economic costs associated with adapting in place. For example, to better understand and anticipate flooding events, a denser sensor network may be required, which would in turn lead to the need for resources for interpreting and making use of the new data. Other costs may include investment in real time stormwater controls to mitigate flooding impacts. However, compared to the cost of constructing new flood protection infrastructure, or the cost of abandoning a coastal city, manifest in rebuilding and in lost opportunity, the economic cost of adapting in place would likely be small.

In addition to the economic costs, the social costs of adapting in place would also likely be much smaller in comparison to coastal city abandonment or new built flood protection. In fact, there is potential for the challenge of increased flooding to actually build social capital (Adger, 2003). For example, citizens of coastal cities can become more educated and involved in understanding, adapting to, and preventing flooding (Tiernan, 2017).

In support of a coastal city's ability to adapt in place to increased flood risk, this dissertation focuses on understanding, predicting, and mitigating coastal flooding through four studies. Each of the four studies will focus on one of these three overarching objectives: to understand, to predict, and to mitgate coastal flooding. In the four studies, emerging technologies, approaches, and datasets are investigated as alternatives to more traditional approaches whose limits are pushed by complex urban coastal environments.

Geographically, this dissertation focuses on two coastal cities in Southeastern Virginia USA: Norfolk and Virginia Beach. It has been estimated that the Hampton Roads Metropolitan Area in Southeastern Virginia, of which Norfolk and Virginia Beach are the two most populous cities, is second only to New Orleans, Louisiana in terms of population centers vulnerable to sea level rise (Fears, 2012). Due to very low relief terrain and faster than average subsidence, this region is experiencing relative sea level rise at a rate faster than the global average. Therefore, it is reasonable to imagine that the conditions that the Norfolk and Virginia Beach are experiencing now, including the increase in flooding even on days without rainfall (Sweet and Park, 2014), are a preview of what many coastal cities will experience in the coming decades. Therefore, the insights and experience gained through research in this area now may be beneficial to other coastal cities as conditions worsen from climate change and sea level rise.

The first study (Chapter 2) focuses on the importance of rainfall measurements close to flood-prone intersections. Rainfall is a major driver of flooding even in coastal settings. Therefore, a spatially and temporally detailed understanding of rainfall is needed to best understand and predict flooding. In reality, however, resource constraints limit the density of rain gauges deployment, and radar-derived products, which provide two-dimensional rainfall, estimation have been shown to have insufficient accuracy. Given that rain gauges cannot be deployed everywhere, the objective of this study was to develop a method to quantify the impact of nearby rain gauges in estimating areal rainfall for watersheds. This method was applied to watersheds which drain to flood-prone street intersections in Virginia Beach.

The second study (Chapter 3) focuses on data-driven prediction of coastal street flooding. One of the major impacts of frequent flooding is on the transportation network of coastal cities. To most effectively reroute traffic to minimize delays and maximize safety, street-scale flood prediction is needed. The complex hydrologic environment of a coastal city including tidal interactions, low-relief terrain, and built storm-water infrastructure make using a more traditional, physically-based predictive model difficult. Data-driven models, on the other hand, have some advantages over physically-based models including the ability to "learn" and mimic the relationship between inputs and outputs instead of attempting to simulate physical processes which by necessity requires simplifying assumptions. This study will demonstrate for the City of Norfolk the use of data-driven models as a step toward street-level flood prediction.

The third and fourth studies (Chapters 4 and 5) focus on the ability of active stormwater controls for mitigating flooding and how their effectiveness will change with sea level rise. The majority of stormwater infrastructure is gravity-driven or controlled passively. Active control of stormwater infrastructure, i.e., the use of actuators such as automated pumps and valves to control the flow of stormwater, has the potential to increase the effectiveness of existing stormwater infrastructure without building new or expanding existing infrastructure (Meneses et al., 2018; Wong and Kerkez, 2018; Kerkez et al., 2016). Chapter 4 describes an open-source software tool that was developed to simulate model predictive control (MPC) of a stormwater system using the Environmental Protection Agency Stormwater Management Model (EPA-SWMM5). A open-source tool implementing MPC for SWMM5 is valuable, but does not currently exist.

As sea levels rise, the designed capacity of existing, gravity-driven stormwater infrastructure will decrease, therefore increasing the importance and utility of active stormwater controls. Chapter 5 provides insight as to how the utility of active stormwater control will change over time given different sea level rise projections and control configurations. This is accomplished by modeling a particularly flood-prone area in Norfolk, the Hague.

This chapter was published in July 2017 in the *Journal of Hydrologic Engineering*.
The final publication is available at https://doi.org/10.1061/(ASCE)HE.1943-5584.0001563

Chapter 2

Effect of rain gauge proximity on rainfall estimation for problematic urban coastal watersheds in Virginia Beach, VA USA

2.1 Introduction

Coastal cities are becoming increasingly vulnerable to flooding (Nicholls and Cazenave, 2010). Recent extreme events, such as hurricanes and tropical storms, have caused severe damage, costing major coastal cities billions of dollars and thousands of lives (Kates et al., 2006; Galarneau et al., 2013). In addition to extreme high return-period events, small return-period rainfall events can also cause flooding in highly urbanized coastal cities and, while less dramatic, these floods can incur significant economic and social costs (Suarez et al., 2005). These lower return-period floods have been occurring more frequently in coastal cities in recent years due to climate change and sea level rise (Ezer and Atkinson, 2014; Sweet et al., 2014). Coastal cities typically have very low topographic relief, large portions of impervious surfaces, a high water table, and tidal influences, which combine to make drainage problematic even without the effects of sea level rise (Titus et al., 1987). Rising sea levels exacerbate drainage problems in coastal cities as tide and groundwater levels rise with the sea level (Bjerklie et al., 2012; Rotzoll and Fletcher, 2012). To add to the flooding problems coastal cities face, precipitation events in general are projected to increase in intensity due to climate change

(Alexander et al., 2006; O'Gorman and Schneider, 2009). To understand and accurately forecast flooding in urban, coastal environments, spatially and temporally detailed rainfall data are needed (Smith et al., 2007), but usually unavailable (Hill, 2015). The typical urban watershed is small in area and has a large proportion of impervious surfaces. This results in a short runoff response time, which increases the risk of flash flooding (Hall, 1984; Fletcher et al., 2013).

Often, neither rain gauge networks nor weather radar can provide rainfall measurements at spatial and temporal resolutions needed to make accurate flood forecasts for urban environments (Hill et al., 2014); rain gauge networks are generally too coarse spatially to provide such detailed information (Seo, 1998). The typical weather radar rainfall product has a spatial resolution of 2 km or coarser (Krajewski and Smith, 2002; Nesbitt and Anders, 2009), which may also be too coarse for urban hydrology applications (Smith et al., 2007). Furthermore, weather radar measurements are indirect requiring an empirically-derived relationship between reflectivity and actual rainfall on the ground (Smith and Krajewski, 1993), and is therefore inherently uncertain to some degree. Efforts have been taken to blend rain gauge and weather radar data (Ercan and Goodall, 2013; Velasco-Forero et al., 2009; Seo, 1998; Sun et al., 2000), but spatially detailed and accurate rainfall data often remain a limiting factor in research and flood forecasting (Hill et al., 2014). To increase the spatial coverage of rainfall estimation, less traditional technologies such as measuring signal attenuation between cell phone towers (Overeem et al., 2013; Zinevich et al., 2008) and using simple, more widespread binary rainfall sensors (Hill, 2015) have recently been evaluated, but are still not widely used and, like radar, do not directly measure rainfall.

Although it is generally accepted that spatially and temporally dense measurements are needed to capture storm events relevant to urban hydrology, the degree of spatial and temporal density required is uncertain. Rainfall spatial variability and its effect on hydrology have been studied using both rain gauge networks (Pedersen et al., 2010; Serinaldi, 2008; Jensen and Pedersen, 2005) and weather radar (Krajewski et al., 2003; Smith et al., 2007). Ciach and Krajewski (2006) used 25 rain gauges stations in a 3 km X 3 km grid to observe small-scale spatial and temporal rainfall variation. In their findings, rainfall exhibited high spatial variability with correlation coefficients decreasing between rain gauges at a 4 km separation distance and a 15-minute time step; the correlation coefficients were lower at a 5 minute time step. Emmanuel et al. (2012) analyzed rainfall radar images finding rainfall patterns to be very spatially heterogeneous with decorrelation distances (the distance at which minimal spatial correlation between two points exists) as low as 5 km. Berne et al. (2004) used geostatistics with rain gauge and an X-Band weather radar data to suggest a simple empirical relationship between watershed area and the corresponding necessary temporal resolution of rainfall observations. They then related the temporal resolution to the needed rain gauge spatial resolution. Their findings suggest that watersheds with areas less than 1 km^2 should have rainfall measurements at about a 3 min temporal resolution and a 2.5 km spatial resolution.

The objective of this paper is to quantify the effect of rain gauge proximity on area-averaged rainfall estimation for small (<1 km^2) problematic urban watersheds in a coastal environment using Virginia Beach, VA USA as a case study. Previous studies, as described above, have tended to focus on describing the temporal and spatial characteristics of rainfall generally. This paper adds to the topic by exploring the ability to estimate rainfall for specific urban watersheds with known flooding problems using available rain gauge networks. This paper also presents a general method to determine how proximity of rain gauges to a watershed impacts the areal averaged rainfall estimates for that watershed. First, the station nearest to each focus watershed is removed from the Kriging interpolation routine to measure the difference in area-average rainfall estimated including and then excluding this nearby station. Cheng et al. (2012) also used a method of systematically removing stations from a Kriging rainfall interpolation procedure, but did not focus on quantifying the impact for estimating rainfall over known problem watersheds. Second, the benefit of rain gauge proximity in areal rainfall estimation is further explored by iteratively removing the next nearest rain gauge to the target watershed in the rainfall estimation procedure. This analysis reveals the benefits of having rain gauges within a given distance of the target watershed. Each scenario is conducted for rainfall observations over three different time scales, 15 minutes, hourly, and daily to understand how these time scales factor into the rainfall estimation for the focus watersheds.

The remainder of the paper is organized as follows. First, the methods and data used to assess the effect of rain gauge proximity on flood prone watersheds in Virginia Beach are described. Second, the results are presented and discussed in terms of the benefit of nearby rainfall observations on watershed rainfall estimation. For context, the results are compared to NEXRAD-derived rainfall estimates as an alternative means for generating area-averaged rainfall estimates. Finally, conclusions and recommendations drawn from this study area are given.

2.2 Methods

2.2.1 Study Area and Focus Watersheds

Virginia Beach is the most populous city (pop. 450,980) in the Commonwealth of Virginia (U.S. Department of Commerce, 2012). The study area is the most populated portion of Virginia Beach. Shown in Figure 2.1, it is 370 km^2, 57% of the total city area, and roughly the northern half of the city.

Specific road intersections with recurrent flooding problems were provided by city engineers and public works division employees. The drainage area corresponding to each of these points was delineated using a 1

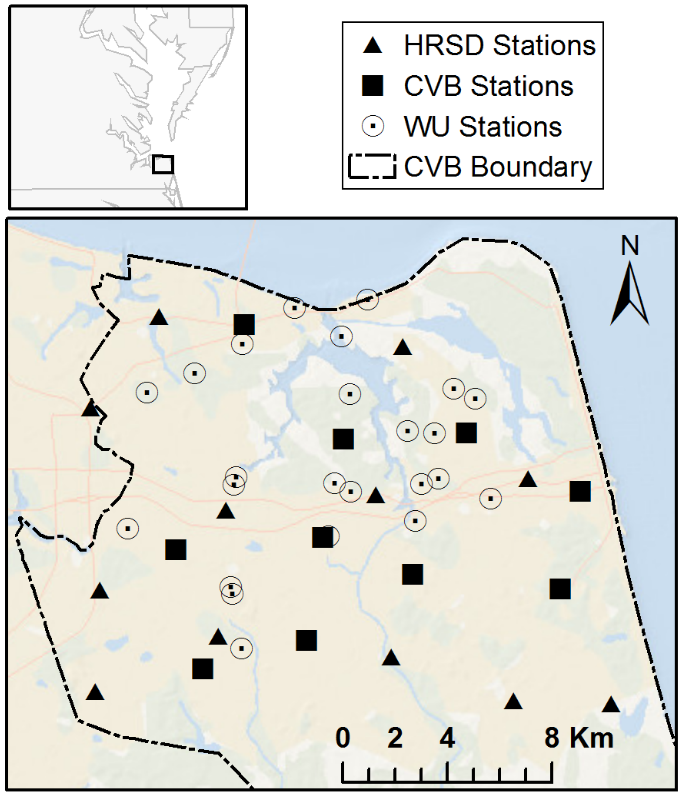

Figure 2.1: Study area, the northern portion of Virginia Beach.

m by 1 m resolution digital elevation model (DEM). These sub-watersheds are shown in Figure 2.2 and their characteristics are given in Table 2.1. For each sub-watershed, the percent imperviousness was obtained from the National Land Cover Dataset 2011 (Homer et al., 2015) and the average slope was calculated from the DEM.

Table 2.1: Focus watershed areas

ID	Description	Area (km^2)	Imperviousness (%)	Ave. Slope (%)
WS-1	Shore Dr & Great Neck Rd	0.76	59	4.7
WS-2	Shore Dr & Red Tide Rd	0.15	69	0.6
WS-3	Ocean View Ave & Mortons Rd	0.02	43	5.3
WS-4	S. Rosemont & S. Plaza Tr	0.13	61	3.9
WS-5	S. Rosemont & Clubhouse	0.26	26	4.4
WS-6	21st & Baltic	0.08	46	3.3
WS-7	Shore Drive & Kendall St	0.69	9	11

2.2.2 Rainfall Data

Rainfall data measured at the Oceana Naval Air Station were used to select focus dates for the analysis. The Oceana Naval Air Station rain gauge was used to select the dates because it is the only weather station run

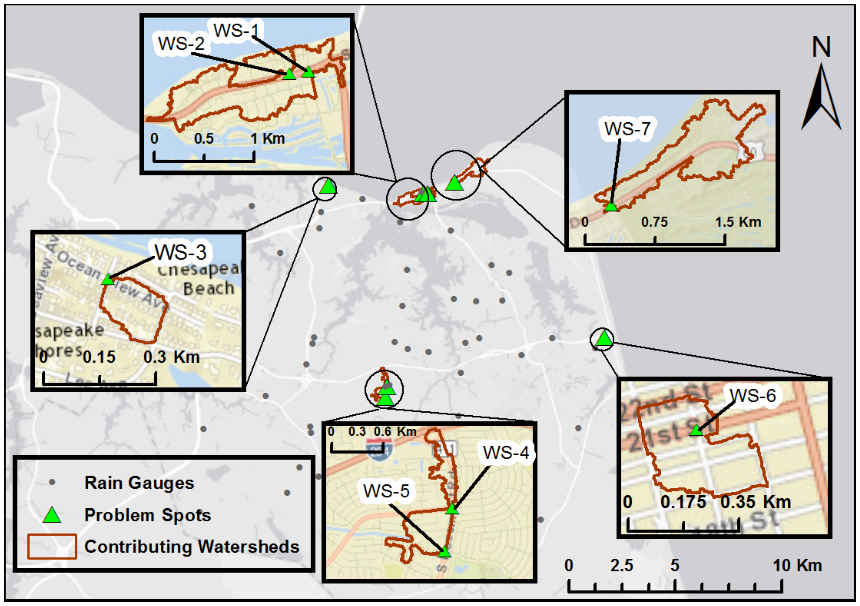

Figure 2.2: Focus watersheds which drain into known flood prone locations.

by the U.S. National Weather Service located within the study area. The 20 days with the highest daily accumulated rainfall at the Oceana Naval Air Station were selected as focus dates for the analysis. For these dates, precipitation data were obtained from three different sources: The City of Virginia Beach, Hampton Roads Sanitation District, and Weather Underground. Figure 2.3 shows the daily rainfall totals averaged over all of the rain gauges from the three data sources for the 20 focus dates. The standard deviations of these averages are also shown to demonstrate the range in rainfall variability in the 20 days analyzed.

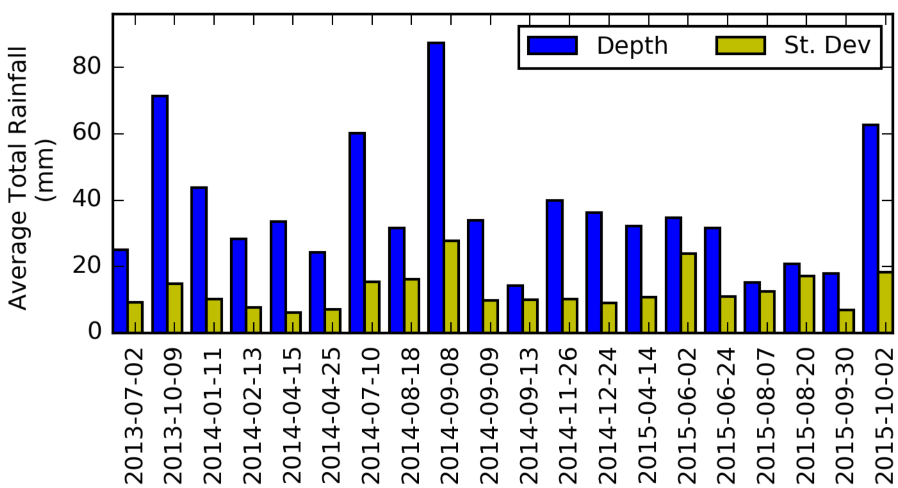

Figure 2.3: Total daily rainfall averaged over all rain gauges for focus dates

The City of Virginia Beach (CVB) in the past five years has installed a rain gauge network consisting of 14 stations. Precipitation data from the 10 gauges within the study area were obtained for the 20 days analyzed. The Hampton Roads Sanitation District (HRSD) has a network of over 50 rain gauges in the Hampton Roads region, 12 of which are within the study region. The data from these 12 gauges, which are quality controlled by the HRSD, were also obtained for the 20 days analyzed. Rainfall data were also obtained from Weather Underground (WU) (http://www.wunderground.com/). WU includes more than 100,000 personal weather stations, purchased and maintained by citizens with their data accessible through the WU site. There were between 7 and 21 WU personal weather stations that reported rainfall values in the study area for the 20 days examined; that corresponds to a 32-95% increase in the number of rain gauges in the study area. Increased spatial coverage has obvious benefits in better understanding spatially heterogeneous precipitation events. However, the data are collected by individual citizen scientists without a quality controlling standard, therefore the validity of the data is uncertain. The process for screening this dataset to identify invalid observations is described in the following section.

Rainfall data from the three sources, CVB, HRSD, and WU, were observed at different temporal resolutions. CVB and HRSD had measurement intervals of 5 and 15 minutes, respectively. Because the WU rain gauges are owned by individuals, the temporal resolution of measurements at these stations varied with an average of 6.2 minutes between observations. Three temporal resolutions were used for the analysis: daily, hourly, and 15 minute. The finest resolution used in the analysis was the coarsest of the three sources: 15 minutes from the HRSD data. The measurements at a temporal resolution finer than 15 minutes were aggregated so that all of the measurements were on a consistent same time scale. The measurements were also aggregated to hourly and daily time scales to study the effect of the time scale on area-averaged rainfall estimation.

2.2.3 Analysis

Quality Controlling of Rainfall Data

The rainfall data from all three data sources were quality controlled first to identify stations not functioning properly on the focus dates. If a station recorded a daily rainfall total of zero for any of the 20 days analyzed, it was assumed that the station was not functioning properly and all values from that station that day, were disregarded. Given that these were the 20 days with the greatest total rainfall over the period of analysis, if a rain gauge did not record any rainfall it was assumed that the rain gauge was not working on that day.

Because the WU data were considered to be less reliable for the reasons described above, these data underwent a simple quality control procedure to identify potentially errant measurements on a 15-minute time scale. At each station, inverse distance weighting (IDW) was used to predict the rainfall based on recorded values from its neighboring CVB and HRSD stations. A minimum of three CVB and HRSD stations, all within 5 km, were used for the estimation. For each 15-minute time step, the IDW estimated value was compared to the value recorded at the WU station in question. If the difference between the predicted value and the recorded value was greater than three times the standard deviation of the measurements used for the IDW prediction, then that measurement was flagged as an outlier. The measurement was also flagged as an outlier if the IDW predicted value was greater than 10 mm and the recorded value was zero. This was to check for stations being off-line intermittently during the focus dates.

The quality control procedure described above was also performed for the CVB and HRSD stations. The number of outliers found from the CVB and HRSD stations served as a baseline against which to compare the number of outliers recorded by the WU stations. If the number of outliers recorded by a WU station was significantly higher than the number recorded by the CVB and HRSD stations, it was judged that the outlying measurements recorded were not the result of real spatial variation, but were the result of measurement errors. The data from the station were therefore disregarded.

Rainfall Interpolation using Kriging

Ordinary Kriging, which assumes a constant, unknown mean over the search neighborhood, was used to quantify the importance of rain gauge proximity in estimating rainfall depth over the seven focus watersheds. It would have been preferable to use Kriging with external drift (Kebaili Bargaoui and Chebbi, 2009); however, this technique requires another related but independent variable such as elevation (Goovaerts, 2000). In this case, since the study area is located near the coastal plane, its elevation is effectively constant. A spherical model was used (Equation 1) for the Kriging semi-variogram models where γ is the variogram, h is the lag, c is the sill, and a is the range. The semi-variograms were assumed to be isotropic. The model parameters (sill and range) were automatically optimized using the RGeostats package (Renard et al., 2015) in R and were calculated for each individual time step with all available data. The Kriging process produced two outputs: the predicted rainfall and the Kriging variance. Kriging variance is the variance of the predicted rainfall value and is a measure of prediction confidence. Both of these outputs were produced as continuous raster datasets and their values were averaged over the area of the focus watersheds to obtain area-averaged estimates of rainfall and Kriging variance using the ArcGIS's Zonal Statistics as Table tool.

$$\gamma(h) = \begin{cases} c\left[\frac{3}{2}\frac{h}{a} - \frac{1}{2}\left(\frac{h}{a}\right)^3\right] & \text{if } h \leq a \\ c & \text{otherwise} \end{cases} \tag{2.1}$$

Experiment 1: Removing the Nearest Stations to a Watershed To quantify the role of nearby rain gauge stations in area-averaged precipitation estimates for the seven focus watersheds, the rainfall was estimated with and then without the nearest quality controlled (from HRSD or CVB) rain gauge station. Any WU rain gauges closer than the nearest quality controlled rain gauge were also removed. For example, considering the watershed diagram in Figure 2.4; S_2 would be excluded, being the closest quality controlled station, and S_1 would also be removed because it is a WU station closer to the watershed centroid than S_2. The number of rain gauges removed for each focus watershed and the average distance of the removed rain gauges are shown in Table 2.2. The table also gives distances for the next nearest station that was not removed in the analysis (i.e., S_3 in Figure 2.4).

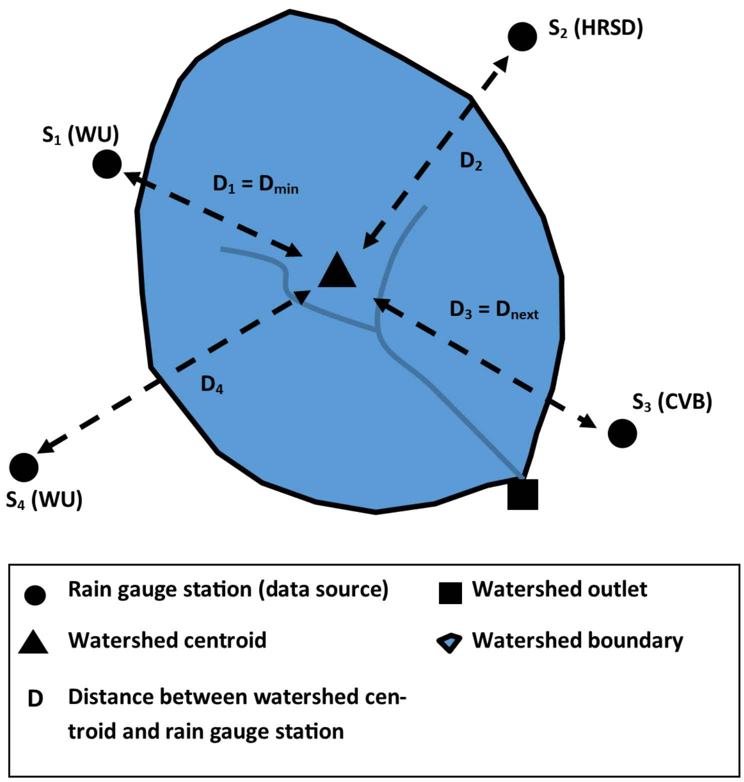

Figure 2.4: Diagram illustrating two Kriging experiments

Table 2.2: Gauges removed in Experiment 1. Distances are measured from watershed centroid.

ID	Num.gauges removed	Distances of removed gauges (m)	Distance to next nearest station (m)
WS-1	3	318, 1419, 2288	2518
WS-2	3	153, 1624, 2340	2643
WS-3	1	1266	1605
WS-4	2	458, 466	1671
WS-5	2	380, 542	2414
WS-6	1	625	2449
WS-7	2	2645, 3016	4201

Two variables were calculated to quantify the impact of removing the nearby stations on rainfall estimates. The first variable is the average increase in variance ($\overline{\Delta \mathrm{Var}_i}$) and the second variable is the average absolute difference in rainfall estimation ($|\overline{\Delta \mathrm{R}_i}|$). These variables were calculated using Equations 2 and 3, respectively. In Equation 2, $\mathrm{Var}_i(t)'$ is calculated the same way as $\mathrm{Var}_i(t)$, the variance of the estimated rainfall on watershed i at time t, but it is calculated without the nearby station(s). Similarly $R_i(t)'$ in Equation 3 is calculated in the same way as $R_i(t)$, the estimated rainfall on watershed i at time t, but without the nearest station(s). As shown in the equations, the difference between $\mathrm{Var}_i(t)'$ and $\mathrm{Var}_i(t)$, and $R_i(t)'$ and $R_i(t)$ is averaged over all values in the time series, $t = 0, ..., N$. Equations 2 and 3 were applied to each of the three time scales considered in the analysis: 15-minute, hourly, and daily.

$$\overline{\Delta \mathrm{Var}_i} = \frac{1}{N} \sum_{t=0}^{N} \frac{\mathrm{Var}_i(t)' - \mathrm{Var}_i(t)}{\mathrm{Var}_i(t)} \tag{2.2}$$

$$|\overline{\Delta \mathrm{R}_i}| = \frac{1}{N} \sum_{t=0}^{N} \left| \frac{R_i(t)' - R_i(t)}{(R_i(t)' + R_i(t))/2} \right| \tag{2.3}$$

Experiment 2: Removing Stations at Increasing Distances from a Watershed A second experiment was performed to quantify the effect of excluding rain gauge stations, at increasing distances from a watershed, on rainfall estimation for that watershed. For each focus watershed, the nearest rain gauge to the watershed centroid was excluded from the rainfall estimation. Unlike Experiment 1, there was no distinction made between stations from different sources in Experiment 2. After the nearest station was excluded, the two nearest stations were excluded from the rainfall estimation. This was repeated until all rain gauges within 8 km of the watershed centroid were excluded. Note that the cases in Experiment 1 are also contained in Experiment 2. Experiment 2, therefore, can be thought of as an extension of Experiment 1. Similar to Experiment 1, in Experiment 2 the average increase in variance ($\overline{\Delta \mathrm{Var}_i}$) and the average absolute difference in rainfall estimation ($|\overline{\Delta \mathrm{R}_i}|$) were calculated for the seven watersheds and the three time scales. However, while in Experiment 1 Equations 2 and 3 were applied using all stations in the study area, in Experiment 2 these equations were applied using only a subset of the stations in the study area, i.e. with removing individual stations up to 8 km from the watershed centroids as explained above.

Comparison with Weather Radar Data

Given that dense rain gauge networks are rare, and that weather radar products are much more widespread (the vast majority of the continental United States is covered by NEXRAD radars), the results of the analysis were compared to weather radar-derived rainfall estimates. The weather radar data used for the comparison was the NEXRAD Level III product "DAA," which is the one-hour precipitation accumulation estimation produced using the Quantitative Precipitation Estimate (QPE) dual-polarization precipitation algorithm which has 256 possible data levels (NOAA, 2017e). This product was selected instead of the one-hour precipitation accumulation calculated using Precipitation Processing System (PPS) based on results from Wu et al. (2012) who found QPE to be generally more reliable than PPS. The spatial resolution of the data was 0.24 km x 1.5 km grid. The NEXRAD data were obtained for each hour of the 20 days analyzed and converted from native binary to a Geotiff raster format using the NOAA Weather and Climate Toolkit. The Zonal Statistics as Table tool in ArcGIS was then used to obtain the average value of the NEXRAD raster cells that intersected each watershed. These area-averaged estimates were compared with the area-averaged estimates obtained from the Kriging of the rain gauge data.

2.3 Results and Discussion

2.3.1 Quality Controlling Results

Results from the quality control showed that 8 of the total 44 stations recorded zero rainfall for at least 12 of the 20 days analyzed (Table 2.3); 6 of these 8 stations were WU stations and 2 were HRSD stations. WU station "KVAVIRGI105" had the most occurrences with 8 of the 20 days reporting zero rainfall. Stations were excluded from the analysis for the day(s) that zero rainfall was reported. It is possible that one rain gauge in such a large area would have had no rainfall even when most rain gauges did. That said, on 8 of the 12 days on which zero rainfall was recorded somewhere in the study area, zero rainfall was only reported at one rain gauge for the day. Since every other rain gauge was reporting at least some rainfall for each of those 8 days, it was assumed that the one that recorded zero rainfall was not functioning. On 3 of the 12 days, 2 stations reported zero rainfall, and on 1 day (2015-09-30) 4 stations recorded zero rainfall. In these cases, since no two stations closest to each other both reported zero rainfall, these stations were also assumed to be not functioning.

The results of the quality control procedure used to identify anomalous measurements from the stations are shown in Figure 2.5. From a total of 31,095 observations, 2,045 were identified as outliers. More than 99% of the outliers identified were three or more standard deviations away from the estimate based on their neighbors, compared to less than 1% which were recorded as zero when the IDW estimate was 10 mm or greater. Eight of the top ten stations in terms of percentage of outliers were WU stations. The WU stations had a higher average percentage of outliers overall (Table 2.4). Only one station stood out statistically: "KVAVIRGI52". The percentage of measurements classified as outliers from this station was 3.6 standard

Table 2.3: Stations and dates where a value of zero rainfall was recorded

StationID	Source	Zero-sum Dates
KVAVIRGI71	WU	2014-09-13
KVAVIRGI79	WU	2014-09-13
KVAVIRGI88	WU	2015-09-30, 2015-10-02
KVAVIRGI105	WU	2014-11-26, 2015-04-14, 2015-06-02, 2015-08-07, 2015-08,02, 2015-09-30, 2015-10-02
KVAVIRGI116	WU	2015-09-30
KVAVIRGI65	WU	2015-09-30
MMPS-036	HRSD	2014-08-18, 2014-09-08, 2014-09-09
MMPS-160	HRSD	2014-04-14

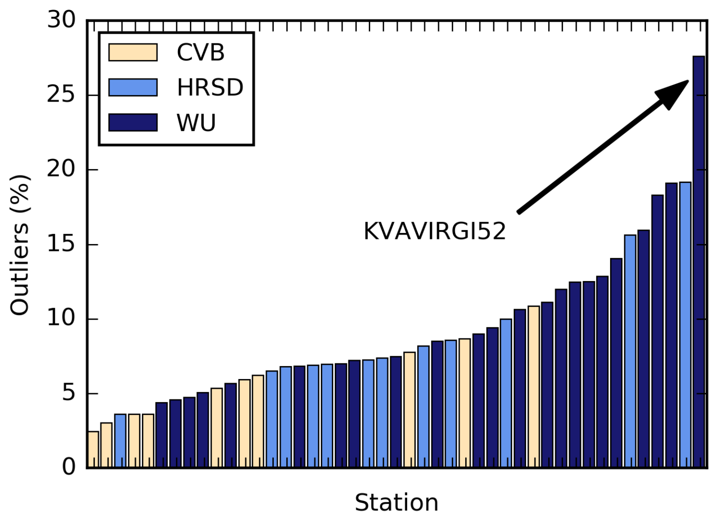

Figure 2.5: Percent of outliers for each station

Table 2.4: Percentage of outliers for each data source

Data Source	Average Percentage of Outliers (%)
CVB	4.5
HRSD	6.9
WU	8.6

deviations from the mean percent number of observations classified as outliers when considering all stations. "KVAVIRGI52" was therefore excluded from the analysis, while all other stations were kept.

2.3.2 Exploratory Analysis Results

Figure 2.6 shows the daily rainfall depths for each station across the study area. There are clear differences in rainfall magnitude and spatial variation between dates. For example, considering the daily total rainfall values for 2014-12-24, 2015-04-14, 2015-06-02, and 2015-06-24, it is clear visually that the daily rainfall on 2015-06-02 is more spatially heterogeneous than the other dates. The data in Table 2.5 confirm this quantitatively. These four daily totals are quite similar: 36.2, 32.6, 34.3, and 31.8 mm, respectively. However, their standard deviations are more variable: 8.9, 11.0, 23.9, and 10.8 mm, respectively. The spatial variation seen visually in Figure 2.6 maybe best be explained quantitatively by the standard deviation to mean ratio, or coefficient of variation (CV). Contrasting the plots of the two dates with the lowest CV, 2014-04-15 (0.18), and the highest CV, 2015-08-20 (0.85), the spatial uniformity on 2014-04-15 and the spatial non-uniformity on 2015-08-20 are clearly seen (Table 2.5).

Although the rainfall can be relatively spatially uniform for a given day, when considering a shorter time step, the spatial variation is often much higher. This can be seen in the 15-minute and hourly CVs which are, on average, at least four times higher than the daily CV (Table 2.6). Even the date with the smallest daily CV, 2014-04-15 (0.18), shows considerable variability in rainfall amounts across stations at the 15-min time step (CV up to 0.65 at 11:15:00; Figure 2.7). This variability on smaller time steps is important to consider for flood forecasting applications in small, highly impervious watersheds with flashy responses. The uncertainty caused by this variation could result in inaccurate predictions about the potential for flooding risk.

Table 2.5: Summary data for daily rainfall

Date	Mean (mm)	Standard Dev. (mm)	CV
2013-07-02	25.1	9.2	0.36
2013-10-09	68.8	19.9	0.29
2014-01-11	43.7	10.1	0.23
2014-02-13	28.0	7.7	0.28
2014-04-15	33.5	6.1	0.18
2014-04-25	24.2	7.1	0.29
2014-07-10	58.5	18.1	0.31
2014-08-18	31.6	16.3	0.51
2014-09-08	84.7	30.9	0.36
2014-09-09	33.1	10.5	0.32
2014-09-13	13.8	10.1	0.73
2014-11-26	39.3	10.6	0.27
2014-12-24	36.2	8.9	0.25
2015-04-14	32.6	11.0	0.34
2015-06-02	34.3	23.9	0.70
2015-06-24	31.8	10.8	0.34
2015-08-07	15.0	12.4	0.82
2015-08-20	20.4	17.3	0.85
2015-09-30	17.8	6.9	0.39
2015-10-02	61.5	19.8	0.32

Table 2.6: Average CV for the time steps examined

Time scale	Average CV
15 minute	2.1
Hourly	2.0
Daily	0.5

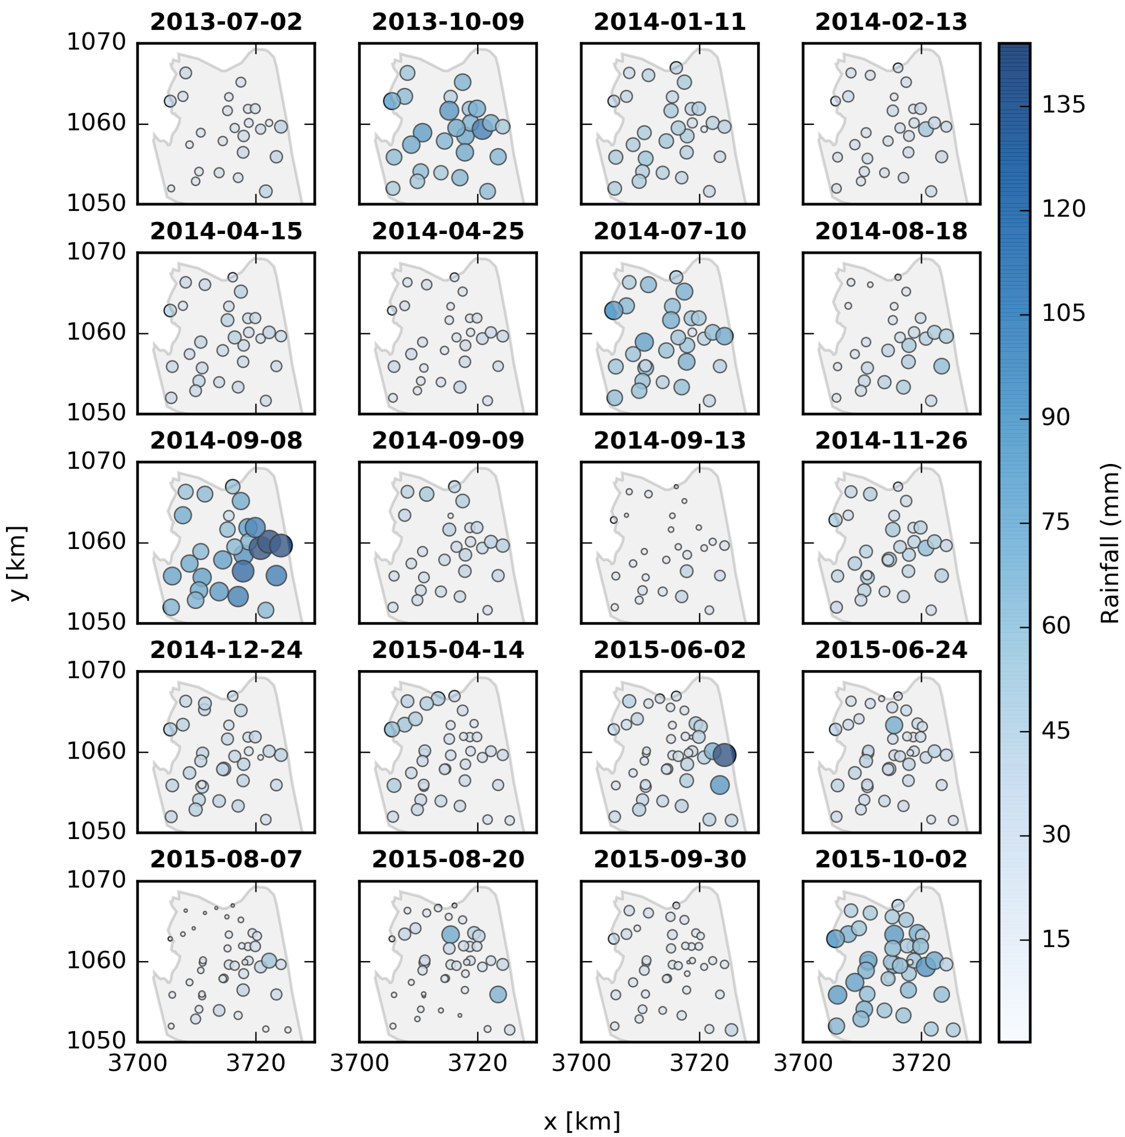

Figure 2.6: Daily rainfall values at each station (size of the scatter points corresponds to relative magnitude of rainfall)

2.3.3 Experiment 1 Results

The results of the rainfall estimation without local information are summarized in Figure 2.8. For most of the watersheds, the variance increased by more than 100% and, generally, the magnitude of the increase corresponded to the distance from the watershed centroid to the nearest excluded rain gauge station (see Table 2.2). For example, the watersheds that had the greatest distance from their centroid to the nearest rain gauge, WS-7 (2645 m) and WS-3 (1266 m), had the smallest increase in variance when those rain gauges

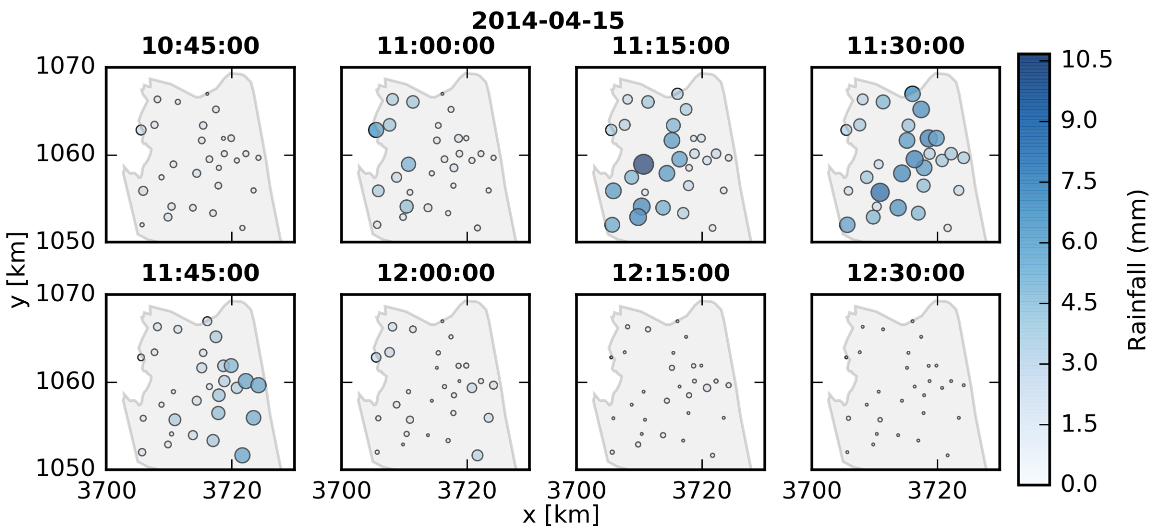

Figure 2.7: 15-min rainfall values for 2014-04-15, a storm with low daily spatial variability (size of the scatter points corresponds to relative magnitude of rainfall).

are removed. Conversely, the watershed that had the smallest distance between its centroid and the closest rain gauge, WS-2 (153 m), had the largest increase in variance when that rain gauge was removed. For the most part, these values vary little between the three time scales (15 min, hourly, and daily). This is because the variance is related more to the spatial arrangement of the observations than their actual magnitude.

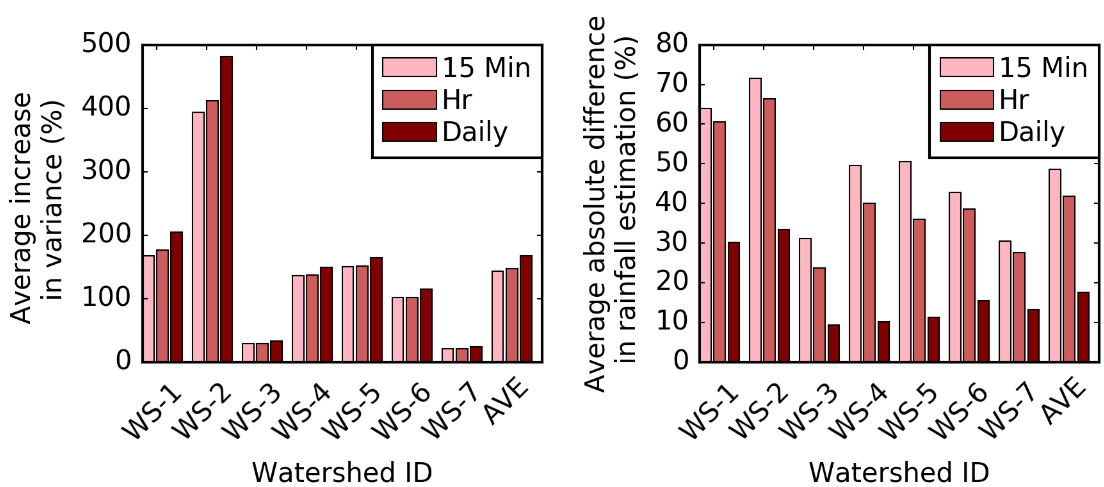

Figure 2.8: Average percent increase in variance and percent difference in rainfall estimation without nearest stations (Experiment 1)

Figure 2.8 shows the average absolute difference in rainfall estimation when the nearby stations are excluded. On average, the percent difference in rainfall estimation for the 15-minute time step was 49% with a maximum of 72% for WS-2. The average absolute difference in rainfall estimation at the 15-minute time

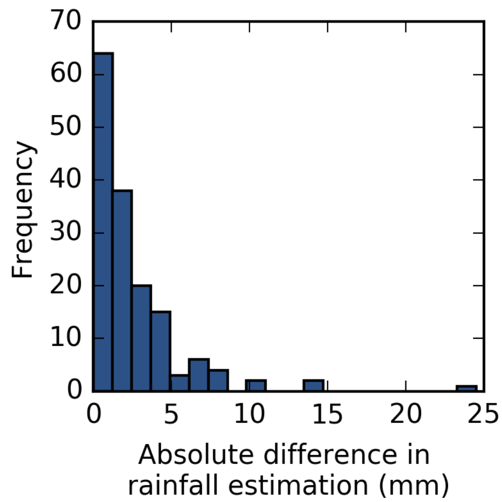

Figure 2.9: Histogram of absolute differences in rainfall estimation in Experiment 1 when rainfall estimates including the nearest gauge were greater than 5 mm

step was 0.34 mm. For 15-minute intervals with more rainfall recorded, the difference was generally larger. Figure 2.9 shows a histogram of the absolute difference in rainfall estimation at the 15-minute time step when the rainfall estimated including the nearest rain gauge was above 5 mm. Out of the 155 data points, 16 had differences greater than 5 mm. The maximum difference in rainfall estimation was 24.5 mm at WS-6, which occurred on 2015-06-02. Without the nearest rain gauge, the estimated rainfall was 7.4 mm; with the nearest rain gauge included, the rainfall estimate was 31.9 mm. For perspective, this difference can be thought of in terms of design storms. The estimated rainfall intensity including the nearby rain gauge (31.9 mm/15 min or 128 mm/hr) for this 15-minute time period corresponded to a 10-year, 15-minute design storm. If this station did not exist, the estimated rainfall intensity (7.4 mm/15 min or 29.6 mm/hr) would erroneously be considered negligible by design standards (Bonnin et al., 2006).

The results from Experiment 1 were used to compare the distances between rain gauges and the focus watersheds, to the rain gauge distance recommendations made by Berne et al. 2004 (Table 2.7). Berne et al. analyzed radar and rain gauge data to recommend temporal and spatial resolution requirements for urban watersheds based on the watershed surface area. Table 2.2 lists the distances of removed rain gauges and Table 2.7 lists the recommended distances for the seven focus watersheds based on their surface area. Four of the seven focus watersheds in this study had more than one rain gauge station within the distance recommended by Berne et al.: WS-1, WS-2, WS-4, WS-5. When the closest rain gauge was removed, the average rainfall estimation changed significantly for each watershed at the 15-minute scale (49%, 60%, 37%, and 44%, respectively). One would expect the difference in rainfall estimation to be much smaller given that one rain gauge station was still within the recommended distance. This suggests that the recommended distances offered by Berne et al. would be too coarse for the study area. The difference in results obtained in this study and Berne et al's may be due to climatic or geographic differences between the study area an urban area on the East Coast of the U.S., and the one used by Berne et al. (2004), a coastal. area in Southern France.

Table 2.7: Experiment 1 results compared to literature-recommended spatial resolutions

ID	Recommended distance to rain gauge (Berne et al., 2004))(km)	Ave. difference at 15 min time scale when nearest rain gauge within the recommended distance is removed
WS-1	2.5	49%
WS-2	2.0	60%
WS-3	1.4	NA
WS-4	1.9	37%
WS-5	2.1	44%
WS-6	1.8	NA
WS-7	2.5	NA

2.3.4 Experiment 2 Results

Results from Experiment 2 show the impact that removing rain gauges at increasing distances from a focus watershed has on rainfall estimation for that focus watershed. Figure 2.10 shows how variance changes when increasing the distance of removed stations from watershed centroids. The change in variance decreases drastically as the distance increases from 0 to 1.5 km, and is effectively negligible by 3.5 km. The greatest change in these results occurs within 1 km. Therefore, on average in this study area, to appreciably increase the confidence of rainfall estimation, a new rain gauge must be within 3.5 km of a given watershed's centroid and would preferably be within 0.5 km.

Figure 2.11 shows how the rainfall estimation changes for each focus watershed as stations are excluded from the Kriging analysis. A linear model was fit to the data for each watershed. For all but WS-4 and WS-5, the R^2 values were at least 0.72, suggesting a linear relationship between distance to the furthest removed rain gauge and the percent change in rainfall estimation for these watersheds.

There may be several reasons why the relationship between distance to the furthest removed rain gauge and the percent change in rainfall estimation for WS-4 and WS-5 appears to be non-linear. One possible explanation is the geography of the watersheds. WS-4 and WS-5 are located more inland compared to the other watersheds. Rainfall variation corresponding to distance from the coast has been previously observed elsewhere (Hayward and Clarke, 1996). The spatial relationship of the neighboring rain gauges may also be a factor. Since WS-4 and WS-5 are located in the center of the rain gauge network, the rain gauges were removed in the analysis from all sides of the watersheds. In contrast, because the other watersheds were on the coast, the rain gauges removed for these watersheds were only from the inland side. Further research specifying the direction from which rain gauges are removed in the analysis would help test this explanation.

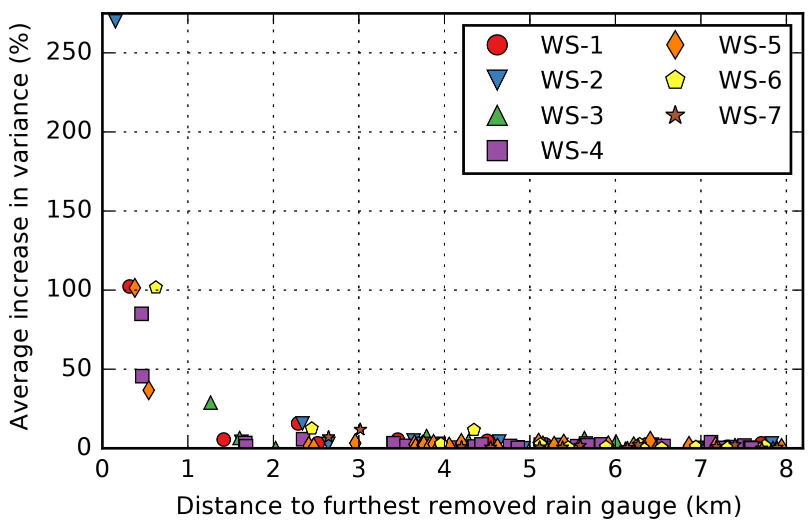

Figure 2.10: Percent increase in variance compared to distance from watershed centroid to excluded measurement stations at 15-minute time scale (Experiment 2)

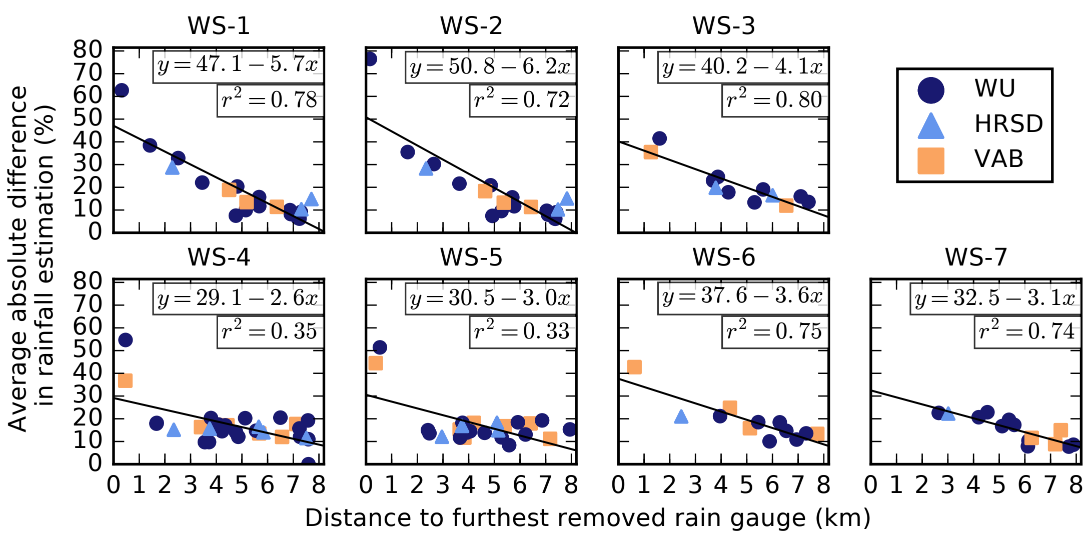

Figure 2.11: Percent difference in rainfall estimations compared to distance from watershed centroid to excluded measurement stations at 15-minute time scale (Experiment 2)

2.3.5 Comparison with Weather Radar Data Results

Figure 2.12 shows the results of the comparison between the area-averaged rainfall estimates based on the Kriging analysis using the rain gauge data, and the area-averaged rainfall estimates from the NEXRAD

DAA data. The figure reports the root mean squared error and Pearson correlation coefficient ("RMSE" and "r" on the figure, respectively). The Figure 2.12a-2.12b show the comparison of the estimates from the two data sources for each of the seven focus watersheds. For these comparisons, the correlation coefficients were between 0.87 and 0.94. Similarly, the RMSE values were between 2.70 mm and 3.71 mm. Many of the hourly rainfall estimates are small (less than 10 mm) and thus would probably not contribute to flooding. To focus on rainfall totals more likely to contribute to flooding, the rainfall estimates greater than 20 mm (based on the rain gauge-based estimate) for all of the focus watersheds were calculated (Figure 2.12h). This subset of the data is much more scattered with a correlation coefficient of 0.72 and a RMSE of 9.72 mm. The increased scatter for large rainfall magnitudes, their coarse temporal (minimum 1 hr) and spatial resolution (1.5 km) limits the ability of NEXRAD DAA rainfall product in urban flash flood applications in this study area.

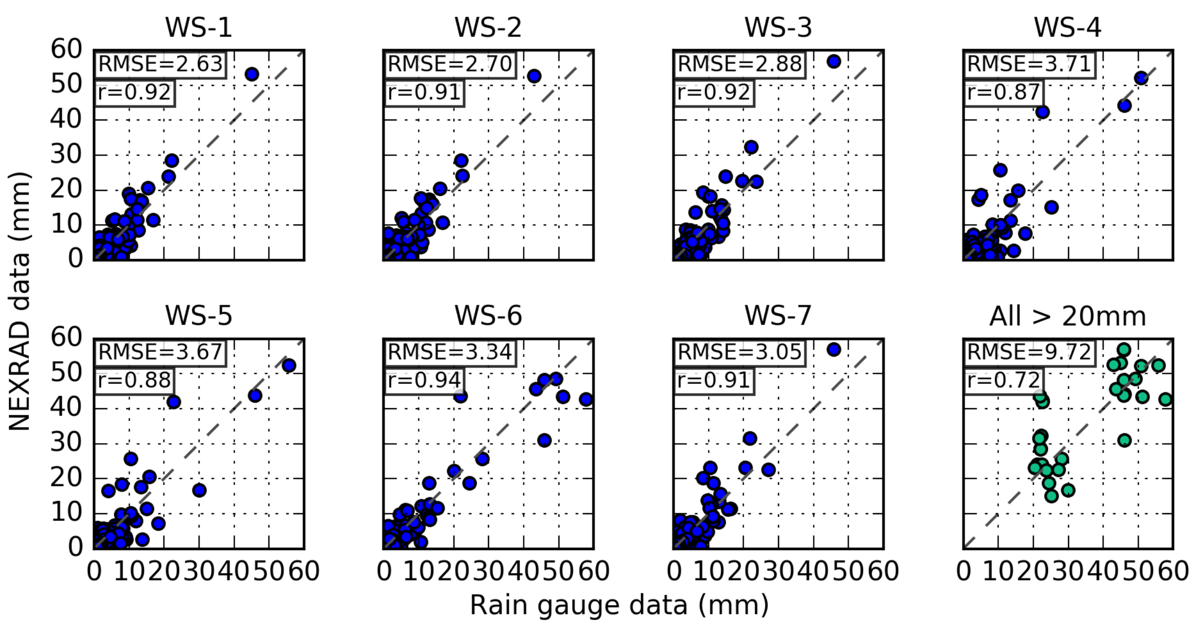

Figure 2.12: Comparison of area-averaged rainfall estimation based on the Kriging of rain gauge data and area-averaged NEXRAD DAA product data. The first seven subplots compare all of estimates by focus watershed. The last subplot (lower right-most) shows the comparison of the rainfall estimates greater than 20 mm (based on the rain gauge estimates) from all focus watersheds.

2.4 Conclusions

The objective of this paper is to quantify the effect of rain gauge proximity on area-averaged rainfall estimation for small (< 1 km^2) problematic urban watersheds. Virginia Beach, VA. served as the case study for the analysis. Rainfall data from three different sources, the City of Virginia Beach (CVB), the Hampton Roads Sanitation District (HRSD), and Weather Underground (WU), were collected. In total, rainfall data from 44 stations for the 20 days with the highest rainfall totals over a three year period were used in the analysis. The WU data were quality controlled on a station by station basis resulting in one station being excluded from the analysis. Kriging was performed to quantify the effect of nearby stations on the rainfall estimation for seven focus sub-watersheds. The results were then compared to radar rainfall estimates for context.

The nearest quality controlled rain gauge to the focus watershed centroids and all closer rain gauges were removed from the network to understand the effect of nearby rain gauges on rainfall estimation. The results of this analysis indicate that rainfall estimations change on average by about 50% across all the watersheds at a 15-minute time step when the nearest station is excluded. For a single watershed, the highest average change in rainfall estimation was over 70% at a 15-minute time step with the largest difference in rainfall estimation of 24.5 mm at a 15-minute time step. This corresponds to the difference between a negligible design storm and a 10-year, 15-minute design storm. Differences of this magnitude could drastically affect flood forecast applications for these small, flashy, urban watersheds.

An analysis was also performed to assess the effect on rainfall estimation and variance from increasingly distant rain gauges to the watershed centroid. The results suggest that rain gauges added within 0.5 km can decrease variance by 50-100% and a rain gauge 3.5 km from the watershed centroid will not decrease estimation variance appreciably. The current rain gauge network has stations within 0.5 km for four of the seven focus watersheds. As flooding problems continue to increase within coastal regions due to climate change and sea level rise, additional problem areas are likely to arise.

To put the analysis of the rain gauge data in the context of radar-derived rainfall products, the rainfall estimation based on a Kriging analysis using the rain gauge data was compared with NEXRAD-derived rainfall estimations. This comparison was on an hourly time step, the finest temporal resolution of the NEXRAD Level III rainfall estimation products. When considering all magnitudes of rainfall estimation, the correlation between the two sources was high (correlation coefficients between 0.87 and 0.94, RMSE between 2.70 mm and 3.71 mm). However, when considering only the higher magnitude estimates more likely to cause flooding (greater than 20 mm) the correlation decreased (correlation coefficient of 0.72, RMSE of 9.72 mm). Besides the decreased correlation with greater rainfall, the unavailability of NEXRAD Level III products at a sub-hourly temporal resolution and coarse spatial resolution of the data (cell sizes of roughly 0.25 km x 1.5 km for the study area) limit current NEXRAD-derived Level III rainfall products in this area for flash-flood applications.

If the long term goal were to cover the entire study area with regularly spaced rain gauges at a 1 km spacing, so that every point would be roughly within 0.5 km of a rain gauge, then 471 stations would be required. Recognizing that this is impractical with existing technology, alternatives could be explored to obtain the dense rainfall data needed to predict flash flooding in an urban setting. An obvious alternative is to focus resources by deploying rain gauges near problem areas. For the current study area, this could begin with adding rain gauges within 0.5 km of the three focus watersheds which do not currently have a rain gauge within that distance. A longer term option may be to refine the use of and advance weather radar technology, specifically focused for fine temporal and spatial resolutions required for urban flash flood warning. Also, advancements in sensor and information technology will also play a role in making denser rainfall networks possible. For example, using acoustic rain gauges, that, unlike the more common tipping bucket rain gauge, do not contain moving parts may reduce costs allowing for larger deployments. Cheaper, binary rainfall sensors could also help fill the need of dense rainfall measurements (Hill, 2015). Leveraging larger efforts like the Internet of Things (IoT) technologies and cyber-physical systems (CPS) approaches, will make it possible to glean meaningful information from observations, and to use this information in stormwater infrastructure controls such as valves, storm gates, and pumps to create a smarter storm water management systems for flood mitigation.

2.5 Acknowledgments

The authors thank the City of Virginia Beach, the Hampton Roads Sanitation District, and Weather Underground for the access to their rainfall data. The authors also acknowledge funding provided by the Mid-Atlantic Transportation Sustainability Center University Transportation Center.

This chapter was published in February 2018 in the *Journal of Hydrology*.
The final publication is available at https://doi.org/10.1016/j.jhydrol.2018.01.044

Chapter 3

Modeling Urban Coastal Flood Severity from Crowd-Sourced Flood Reports Using Poisson Regression and Random Forest

3.1 Introduction

Flooding in low-lying, coastal cities has become more common in recent years due to climate change and relative sea level rise (Sweet and Park, 2014). Relative sea level is expected to rise substantially (Vermeer and Rahmstorf, 2009; Church et al., 2001) will worsen the problem of flooding in coastal cities. Flooding in coastal cities is caused by large, life-threatening, high-return period events such as Hurricanes Harvey and Irma whose flooding recently severely affected coastal cities in Texas and Florida USA, respectively. Additionally, many coastal cities have low-relief terrain and low elevation making stormwater drainage problematic. This can make coastal cities susceptible to flooding from smaller, low-return period events such as severe thunderstorms. The long-term effects of legacy engineering decisions can further add to an urban city's flood risk (e.g., the use of non-engineered fill used to reclaim streams which causes higher than average subsidence rates (Turner, 2004)).

The ability to accurately predict flooding allows decision makers to proactively mitigate the effects of flooding (Zevenbergen et al., 2008), which is key to a city's resilience to natural hazards (Godschalk, 2003).

Accurate flood prediction allows decision makers to maximize safety in the case of large events, and minimize infrastructure damage and social and economic disruption in the case of smaller events. Accurate flood prediction also allows cyber-physical (or smart) stormwater systems to perform optimally, further mitigating the effects of flooding (Kerkez et al., 2016).

Modeling and predicting flooding in urban coastal environments can be challenging. Urban coastal floods are influenced by a combination of different environmental, geographic, and human-related factors (Gallien et al., 2014). Environmental factors that contribute to coastal flooding include rainfall, wind, tide levels, and ground water table levels. Geographic factors such as elevation, soil properties, proximity to the coast, and the land use and land cover of the drainage area can influence whether a given location experiences flooding. In urban settings human-related factors including built stormwater infrastructure and the condition of that infrastructure, which is often underground, also play a role in the location and severity of flooding. For example, clogged stormwater inlets and undersized stormwater pipes can increase the chance and severity of flooding. High tidal levels can inundate stormwater outfalls rendering them ineffective at draining stormwater to the ocean, a condition which will become more frequent with sea level rise. The need to accurately represent such systems and their changing conditions further adds to the complexity of urban flood modeling.

Urban coastal flood events can be modeled using physically-based 1D (Mark et al., 2004a) or 2D models (Mignot et al., 2006; Hunter et al., 2008; Bates et al., 2005; Smith et al., 2011; Gallien et al., 2014). However, the simplified representations of reality used in physically-based models can be a limitation given the combination of variables and their interactions, and the complexity of the physical environment. Two-dimensional, hydrodynamic models make fewer simplifications compared to 1D models, however, this comes at a larger computational cost (Leandro et al., 2009) which makes executing, and especially calibrating, a 2D model difficult (Caviedes-Voullième et al., 2012).

Another modeling approach shown to be effective in many fields (Yang et al., 2017a) including hydrology (Solomatine and Ostfeld, 2008) is data-driven modeling. Data-driven models detect patterns in the data to map model inputs to model outputs without attempting to simulate the physical processes (Solomatine and Ostfeld, 2008). Thus, the relationship between the inputs and outputs is not assumed, as in physically-based models, but learned. While physical processes are not directly simulated using data-driven models, understanding of physical processes usually influences the selection of input variables used to predict the output variable (Booker and Woods, 2014).

The recent increase in availability of earth observation data, coupled with advances in machine learning algorithms, have expanded the possibilities and use of data-driven modeling in hydrology. Machine learning algorithms have been used extensively in hydrology for applications such as predicting reservoir operations (Yang et al., 2016), soil mineral weathering (Povak et al., 2014), streamflow (Yang et al., 2017b; Solomatine and Xue, 2004; Wang et al., 2009), groundwater potential (Naghibi et al., 2017), and groundwater level (Sahoo et al., 2017). Data-driven and machine learning algorithms in flooding applications specifically have been used by Tehrany et al. (2013), Wang et al. (2015), and Tien Bui et al. (2016) who predicted areas susceptible to flooding, Adamovic et al. (2016), who modeled flash flooding on a regional scale, and Solomatine and Xue (2004), who predicted streamflow for flood forecasting. Despite the expanded use of data-driven models in hydrology, few studies have used data-driven methods to model flooding within coastal urban environments. The closest work may be the statistical analysis of tidal records in the United States to estimate the amount of time that coastal cities have experienced flooding in the past several decades and project flooding in the coming decades (Ezer and Atkinson, 2014; Sweet and Park, 2014; Moftakhari et al., 2015; Ray and Foster, 2016).

The objective of this study is to use data-driven modeling to predict flooding severity for a given storm in an urban coastal setting. Crowd-sourced flood reports recorded during flood events are used for model training and are considered a proxy variable for flood severity. Although a more objective measure of flood severity is preferred to the number of flood reports (e.g., flood inundation depth and duration throughout the study domain), often such data is not available. Relevant environmental data (rainfall, tide levels, water table level, wind speed and direction) will be used as inputs to the model.

A data-driven approach is appropriate for this application due to the complexity of modeling urban coastal flooding, as discussed above, which makes using a physical model difficult. This paper will investigate and compare two different data-driven models, Poisson regression and Random Forest regression. Poisson regression is a generalized linear model and was selected because it is commonly used to model rare events (D'Unger et al., 1998) and a flood report, while increasing in occurrence, can still be considered a rare event. Random Forest was selected due to its wide use as a machine learning algorithm in hydrology applications (Yang et al., 2016; Wang et al., 2015; Loos and Elsenbeer, 2011) and other fields (e.g., Mutanga et al., 2012; Svetnik et al., 2003).

The data-driven approach will be applied in Norfolk, Virginia USA. Norfolk and the surrounding Hampton Roads region is one of the most vulnerable metropolitan centers to coastal flooding in the USA (Fears, 2012). Since 2010, the City has collected quality-controlled, crowd-sourced street flooding reports for 45 storm events. In this study, the two data-driven models, Poisson regression and Random Forest regression, will be trained to predict the number of street flood reports per storm event, given the rainfall, tidal, water table, and wind characteristics of the storm event. The models will be evaluated and compared using primarily the root mean squared error (RMSE) and mean absolute error (MAE) between the predicted number of street flood reports and the actual number of street flood reports.

This is a first step toward the use of data-driven approaches in urban coastal flood modeling. Additionally, although Gaitan et al. (2016) employed exploratory methods to glean information from open spatial data, weather data, and user reports, the use of crowd-sourced data in the training and evaluation of data-driven predictive models for urban flood modeling has not been demonstrated or discussed thoroughly in the literature. This is relevant currently as multiple platforms now exist for collecting crowd-sourced information regarding urban flooding (Le Coz et al., 2016) and it can be expected that, due to the nearly universal use of internet connected devices, crowd-sourced data will continue to grow in volume. It is also anticipated that the results of the model will shed light on the relative importance of different environmental factors in predicting coastal flooding, another subject that has been given little attention in previous literature regarding urban coastal flooding.

The remainder of this paper will proceed as follows. First, background will be given describing the study area, the model input and output data, and an introduction to the data-driven models used. Next, the methods are presented describing the preparation of the data for the models and how the data-driven models were applied and evaluated. The model results are then presented and discussed, and finally conclusions are given.

3.2 Study Area, Data, and Model Background

3.2.1 Study Area and Street Flooding Record

Norfolk, Virginia USA, shown in Figure 3.1, is an ideal study area for this research considering its vulnerability to flooding, its economic and military importance, and the availability of quality-controlled crowd-sourced data regarding flood occurrences for the city. Norfolk is one of the most vulnerable cities to coastal flooding in the USA due largely to land subsidence rates causing Norfolk and the surrounding area to experience relative sea level rise at a rate faster than the global average (Kleinosky et al., 2006). As home to the largest terminal of the Port of Virginia, the 3rd most used port on the East Coast of the US (The Port of Virginia, 2016), Norfolk plays an important role in the economy of Virginia and the surrounding states. The world's largest naval base, Naval Station Norfolk, is also within Norfolk, making the flooding risks in the area important to US national security (Broder, 2009).

An important factor in selecting the study area was the availability of crowd-sourced flood record data. A record of reported flooded street locations has been kept in Norfolk starting with Hurricane Nicole on 30 September 2010, shown in Figure 3.1. This is a unique dataset because often, observational data from flooding events is a limiting factor in creating useful flood models (Smith et al., 2011). Because such data is often sparse, photographs of flooded locations and personal interviews have been used out of necessity in the calibration and verification of flood models (Smith et al., 2011). Even satellite imagery has been used to estimate flooding extents (Ireland et al., 2015), but is less useful as a street-scale flood record in an urban setting due to its coarse spatial resolution.

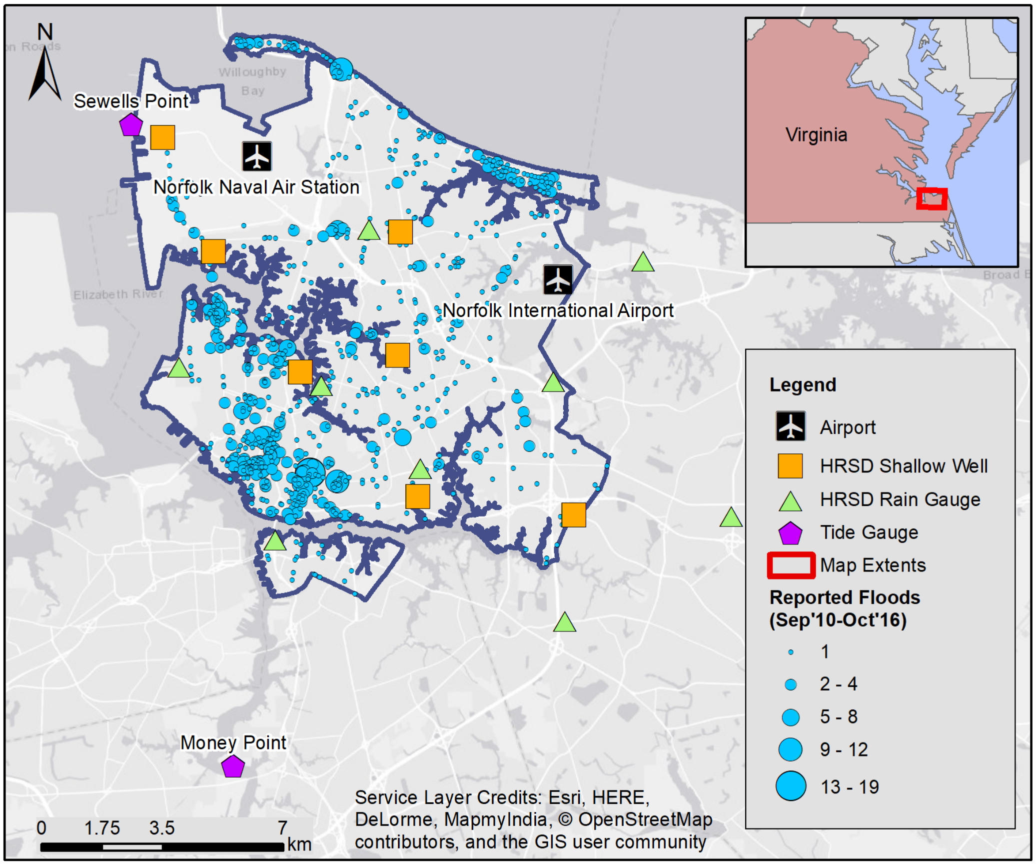

Figure 3.1: Study area: Norfolk, Virginia USA

3.2.2 Description of Model Input and Output Data

The objective of this study is to develop a model capable of predicting flooding severity resulting from a given storm event based on the environmental conditions of that event. The environmental condition input data for the model consisted of rainfall, water table level, wind, and tide level observations. These were obtained from the Hampton Roads Sanitation District (HRSD) and the US National Oceanic and Atmospheric Administration (NOAA). From HRSD, rainfall, water table elevation, and wind direction and wind speed data were obtained. The rainfall observations were on a 15-minute time scale, and the water table elevations and wind data were on a 2-minute time scale. From NOAA, 6-minute water elevations and daily high and low tides recorded at the Sewells Point (NOAA, 2017e) and Money Point (NOAA, 2017c) tide gauges were obtained. Wind speed, wind gust, and wind direction data recorded at the Money Point station at 6-minute time intervals were used as well. Daily rainfall and wind data collected at two airports in the study area, Norfolk International Airport (NOAA, 2017a) and Norfolk Naval Air Station (NOAA, 2017b), were also obtained from NOAA. The rain gauge, water table, wind, and tide gauge stations and airports are shown in Figure 3.1. All of the raw data together consisted of more than 15 million observations. To keep the time series data organized, a simplified version of the Consortium of Universities for the Advancement of Hydrologic Sciences Incorporated (CUAHSI) Observations Data Model (Horsburgh et al., 2009) was implemented in a sqlite database.

The target data used for the model training and evaluation were the number of crowd-sourced flooded location reports per storm event from September 2010 to October 2016. The flood reports were made and catalogued using the City's custom System to Track, Organize, Record, and Map (STORM). STORM is used by the City to record, and catalogue impacts from storms on the City's infrastructure (e.g., downed powerlines, damaged trees) and STORM data are viewable online at http://gisapp1.norfolk.gov/stormmap. For most of the study period, only City of Norfolk staff were able to make reports in STORM however, in the Spring of 2016 the mobile application used for reporting was made available to the public. Reports made by the general public underwent an approval process by City staff.

The two categories of STORM reports used for model training in this study were "flooded street" and "flooded underpass". A total of 45 storm events (listed in Table 3.1) were reported to have caused at least one flooded street or flooded underpass in the period of record. The number of reported floods per event ranged from 1 to 159. Figure 3.2 shows a box plot of the street flood reports per event. Eight of the events were hurricanes; the rest were unnamed or given generic names by city workers. The number of flood reports made from the top six storm events were much larger than the number of reports made from the other 39 storm events. These larger events are marked as points and labeled in Figure 3.2.

3.2.3 Model Alternatives

Poisson Regression

Poisson regression is a generalized linear model (GLM) commonly used to model rare event, count data. Applications of Poisson regression include modeling crime rate (Osgood, 2000), disease incidence (Frome and Checkoway, 1985), and manufacturing defects (Lambert, 1992). Morrison and Smith (2002) and Viglione et al. (2014) used Poisson distributions to model the arrival time and occurrence of flood peaks, respectively.

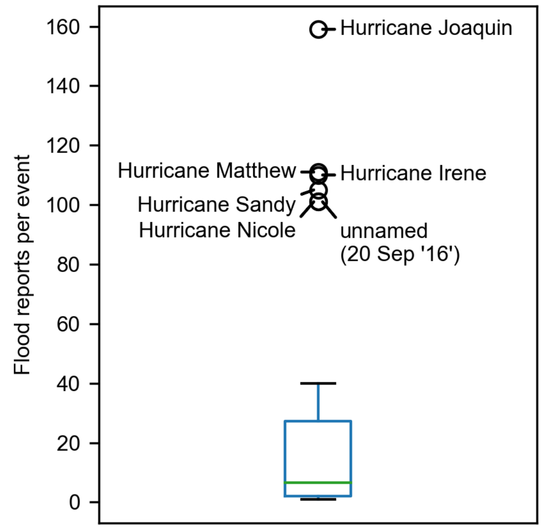

Figure 3.2: Summary plot of reported floods per event in Norfolk, VA Sep. 2010 - Oct. 2016

There are two main assumptions made when using Poisson regression. The first is that the response variable (number of flood reports in this case) follows a Poisson distribution

$$P = e^{-\lambda}\frac{\lambda^k}{k!} \tag{3.1}$$

where P is the probability that k number of events will occur per interval of time and λ is the event rate. The second major assumption when using Poisson regression is that the variance and the mean of the response variable are equal. Thus, the probability distribution (eq. 1) can be specified by only one parameter, λ (Coxe et al., 2009).

In Poisson regression, the mean parameter, λ, is defined by the log-linear function

$$\lambda = e^{-x_i\beta} \tag{3.2}$$

where x_i is a vector of input values for time i and β is a corresponding vector of model parameters, which is optimized during training (Cameron and Trivedi, 1998).

Random Forest

Random forest, developed by Breiman (2001), is an ensemble machine learning algorithm which uses a large number of classification or regression trees (CART) to make a prediction (Breiman et al., 1984). The response variable in this case, the number of flood reports per event, is modeled using regression, therefore the Random

Table 3.1: Events recorded to have caused flooding in Norfolk Sep. 2010 - Oct. 2016

Event Date	Event Name	Flood Reports
29 Sep 2015	Hurricane Joaquin	159
05 Oct 2016	Hurricane Matthew	111
27 Aug 2011	Hurricane Irene	110
28 Oct 2012	Hurricane Sandy	105
20 Sep 2016	unnamed	101
30 Sep 2010	Hurricane Nicole	101
02 Sep 2016	Hurricane Hermine	40
10 Jul 2014	Thunderstorms	39
09 Oct 2013	Heavy Rain	36
16 May 2014	Heavy Rain	35
08 Sep 2014	Rainy Monday	31
20 Jan 2016	January Winter Weather	26
24 Jul 2014	unnamed	18
24 Sep 2015	Noreaster	16
19 Sep 2016	Heavy Rain	11
02 Mar 2015	unnamed	10
11 Jul 2015	Thunderstorm	10
19 Jul 2016	Thunderstorm	9
25 Feb 2016	unnamed	8
03 Jul 2014	Hurricane Arthur	8
31 Jul 2016	Thunderstorm	8
02 Jul 2015	unnamed	7
15 Jan 2016	unnamed	6
03 Jun 2016	Severe Weather - 6/5	6
04 Sep 2014	Thunderstorm	5
19 Jun 2014	Thunderstorms	5
01 Feb 2016	unnamed	5
23 Nov 2014	unnamed	4
13 Sep 2014	Saturday Storm	3
30 Dec 2015	Heavy Rainfall	3
09 Jul 2014	Thunderstorms	2
25 Jul 2016	Bernie (Training)	2
10 Jun 2016	unnamed	2
29 Sep 2014	unnamed	2
16 Dec 2010	Snow	2
24 Feb 2016	February 24th Storm	1
17 Nov 2014	Storm	1
30 Oct 2015	unnamed	1
20 Jul 2016	unnamed	1
17 Sep 2015	unnamed	1
02 Sep 2015	unnamed	1
18 Aug 2014	unnamed	1
24 Sep 2014	Heavy Rain	1
09 Jun 2014	unnamed	1

Forest model is an ensemble of regression trees. In the training of a regression tree, rules based on the response variable are developed to divide observations until the resulting predictions have a minimum amount of node impurity. Node impurity for regression trees, as defined by Breiman et al. (1984), is the sum of the squared deviations between the predicted and actual value (Loh, 2011). The regression tree's rules are a collection of linear divisions of the observation data that, together, create a non-linear decision surface.

One of the main problems of regression trees is that they are prone to overfitting to the training data and thus perform poorly when given unseen data (Murphy, 2012). Random Forest is an approach that attempts to address this weakness. When an individual regression tree is trained in the Random Forest algorithm, a portion of the input records and predictor variables are randomly selected as input to the training. This process is repeated for the number of regression trees specified by the modeler, thus creating a group of regression trees, each trained on a randomly selected subset of the records and input variables. This group of regression trees constitutes a Random Forest model. The prediction made by a Random Forest regression model is the average of the predictions made by each individual regression tree. The random selection of input records and variables in the training of the individual regression trees creates variety in the weak learners, thus avoiding overfitting of the model to the training data.

Beyond the actual predictive capabilities of Random Forest, the algorithm can be used to understand variable importance. Because many regression trees are being produced with different sets of input variables, the Random Forest algorithm learns and records the relative importance of the input variables in predicting the output. This capability is especially attractive as one of the objectives of this study is to understand the relative importance of explanatory variables in predicting street flooding, thus directing future investments in improving observational networks within the city.

3.3 Methods

3.3.1 Input Data Pre-processing

All of the raw input environmental data were aggregated to match the time scale of the flood reports. For all of the days on which no flood reports were made, and for storm events resulting in flood reports made only on one day, the data were aggregated to a daily time scale. For the storm events whose flood reports spanned multiple days, the data were aggregated across the days so that each storm event had only one set of average environmental conditions. For example, flood reports labeled "Hurricane Sandy" were recorded over three days, 2012-10-27, 2012-10-28, and 2012-10-29. The higher high tide taken for this event was the highest of the higher high tides of these three days, the average level of the surficial groundwater table was the average over these three days, and the total cumulative rainfall was the accumulated rainfall from the three days. The resulting dataset consisted of 2,171 records of average environmental conditions, mostly at a daily time scale, from September 2010 through October 2016. The aggregated environmental input variables are shown in Table 3.2.

Different approaches of aggregation were taken for the various environmental data. Four derivatives of the raw HRSD 15-minute rainfall data were included in the models as inputs: total cumulative rainfall, maximum hourly rainfall, maximum 15-minute rainfall, and cumulative rainfall in the previous three days. The different derivatives of the rainfall data were included to account for different types of storm events that may cause flooding. For example, during convective thunderstorms in the summertime, the maximum 15-minute rainfall would be high, but the total cumulative rainfall may be low. For nor'easters, the cumulative rainfall would be high while the maximum 15-minute rainfall may be low.

As with the 15-minute rainfall data, several tide-related variables were model inputs including high and low tides, and average tide level. In coastal cities, such as Norfolk, the timing of rainfall and the tide levels can have an effect on flooding. For example, if tide level is especially high when a large amount of rain falls, the stormwater oulets may be submerged. Such tailwater conditions do not provide sufficient head difference for gravity-driven storm drainage systems to function properly resulting in more flooding than if the tide were low and the same amount of rain fell. To account for such interactions between tide and rainfall, the tide level at the time of the maximum 15-minute rainfall and the tide level at the time of the maximum hourly rainfall were included as model inputs.

The environmental conditions data were averaged across all the stations that recorded the variable. For example, the "Daily cumulative rainfall" is the total cumulative rainfall averaged across all 11 rain gauges. This spatial averaging was done because for some of the stations there was a considerable amount of missing data over the six years of the study period. If the variables were not averaged across measuring stations, it would appear that the stations that had less missing data were more important which would make it more difficult to understand the importance of the actual environmental variables compared to the consistency of measurements at an individual station.

To reduce noise in the data, days on which little or no rainfall was recorded were not used in the modeling procedure. Of the 45 events for which flooding was reported, 42 had an average cumulative rainfall total of 0.25 mm or greater. Of the three events with less than 0.25 mm of rainfall, only one flooded location was reported for two of the events and only two flooded locations were reported for the third event. Given that very minor flooding was reported for days without rainfall, days with little to no rainfall (¡0.25 mm of cumulative rainfall) were considered in the model training and evaluation. This reduced the number of total records used to train and evaluate the model from 2,171 to 814.

Table 3.2: Input feature names and descriptions

Input Feature	Units	Source Organization	Abbreviation
Total cumulative rainfall	mm	airports, HRSD	RT
Maximum hourly rainfall	mm	HRSD	RHRMX
Maximum 15-minute rainfall	mm	HRSD	R15MX
Cumulative rainfall in previous three days	mm	HRSD	R3D
Average water table elevation	m above NAVD88	HRSD	GW_AV
Average tide level	m above MSL	NOAA	TD_AV
Tide level at time of maximum 15-minute rainfall	m above MSL	NOAA	TD_R15
Tide level at time of maximum hourly rainfall	m above MSL	NOAA	TD_RHR
High tide	m above MSL	NOAA	HT
Higher high tide	m above MSL	NOAA	HHT
Low tide	m above MSL	NOAA	LT
Lower low tide	m above MSL	NOAA	LLT
Average daily wind speed	km per hour	airports, HRSD, NOAA	AWND
Average daily wind direction	degrees	airports, HRSD, NOAA	AWDR
Average wind speed over 6-minutes	km per hour	airports, HRSD, NOAA	WSF6
Average wind direction over 6-minutes	degrees	airports, HRSD, NOAA	WDF6
Average maximum 2-minute wind gust over 6-minutes	km per hour	airports, HRSD, NOAA	WGF6
Average wind speed over 2-minutes	km per hour	airports, HRSD, NOAA	WSF2
Average wind direction over 2-minutes	degrees	airports, HRSD, NOAA	WDF2

3.3.2 Model Training and Evaluation

Model training and evaluation were performed using two independent, randomly selected partitions of the output and corresponding input data. In some studies, the dataset is split into three partitions, a training, evaluation, and validation set (Tao et al., 2017), however, since a two-way split is common in this field (Tien Bui et al., 2016; Tehrany et al., 2013; Solomatine and Xue, 2004) and the available dataset is of limited volume, the dataset was split into only two partitions. In the model training, the evaluation dataset was withheld and the models were fit to only the training data. By withholding the evaluation dataset in model training, the models can be evaluated using data not previously seen by the models, thus simulating actual use of the predictive models.

The R programming language (version 3.3.3) was used to partition the datasets, train the two models, and apply the models to the unseen, evaluation dataset (R Core Team, 2017). The dataset of environmental conditions (the model input data) for storm events from September 2010 to October 2016 and the number of reported flood locations for each event (the model output data) were randomly divided into a training set (70%) and an evaluation set (30%) (Tien Bui et al., 2016; Tehrany et al., 2013; Solomatine and Xue, 2004) using the “caret” package in R (Kuhn et al., 2016). This package supports the stratified sampling of the datasets based on the distribution of the model output data, which included the 42 storm events for which flooding was reported and the 772 events for which no flooding was reported. By using stratified sampling, the distribution of the number of reported floods (the output variable) in both the training and evaluation datasets was similar to the distribution of the number of reported floods of the entire data set. To account for potential bias in the division of the data into training and evaluation sets, the random division was made and the models were trained 100 times independently for both models.

Since Poisson regression is a parametric model, the training of the Poisson regression model consisted of optimizing the model coefficients. The built-in "stats" package in R was used for the Poisson regression model. The training of the Random Forest consisted of training each of the individual regression trees making up the Random Forest. The 'randomForest' package (version 4.6.12) was used for the Random Forest model (Liaw and Wiener, 2002).

The Random Forest model has two main parameters, the number of trees per forest, and the number of random predictor variables each tree uses. A sensitivity study of these parameters was performed to determine their effect on model performance. To determine the model sensitivity to the number of trees per forest, the models were trained with the number of trees varying between 2 and 2,000 with the default number of variables per tree (i.e., one-third of the variables, or six in this case). The "tuneRF" function in the "randomForest" package was used to determine the optimum number of variables per regression tree in the Random Forest. This function changes the number of variables used in each regression tree to find the number of variables that minimizes the out-of-bag error within the Random Forest. The out-of-bag error is the prediction error when each input record is applied only to the portion of the regression trees which did not contain that input record in its training sample (Breiman, 2001). Only the training data was used to determine the appropriate number of trees and variables per tree (Xu et al., 2017).

Once trained, both the input training dataset and the input evaluation dataset, which was withheld in the model training, were used as input for the models. Applying the models to the input data produced a predicted number of reported floods for each array of input values. The predicted numbers of reported floods were compared with the known number of reported floods. Root mean squared error (RMSE) and mean absolute error (MAE) between the known and predicted number of flood reports were the two main metrics used to evaluate the models. Since each model was trained 100 times (once for each of the 100 random divisions into training and evaluation data), a distribution of predicted number of flood reports was produced in the model evaluation for each known number of reported floods. To describe these distributions, their standard deviations (std) were plotted and the average standard deviation of each model was reported.

In addition to RMSE, MAE, and std, the models' false negative and false positive predictions were also used to evaluate the models. False negative predictions occur when the predicted number flood reports is zero and the true number of flood reports is non-zero. False positive predictions occur when the true number of flood reports is zero and the predicted number of flood reports is non-zero. These terms are sometime referred to as Type I and Type II errors, respectively (Beguería, 2006)

False negative and false positive predictions are of particular interest to a decision maker. False negative predictions may jeopardize human safety and incur higher costs in recovery when a true positive prediction would have led to less costly, preventative measures. False positive predictions over time can erode trust in the warnings (Basha et al., 2008). For the Random Forest and Poisson regression models, the statistical characteristics (e.g., count, mean, standard deviation) of the false negative and false positive predictions were reported and compared. Since predictions were on a continuous scale and true flood reports were integers, the predictions were rounded to the nearest integer when calculating the false negative and false positive rates. For example, a prediction was considered false positive when the number of flood predictions was at least 0.5 (which would round to 1) and the true number of flood reports was zero.

3.4 Results and Discussion

3.4.1 Model Results

The results of the Poisson regression training and evaluation are shown in Table 3.3 and Figure 3.3. Predictions from the Poisson regression greater than 159 flood reports (the largest number of flood locations reported from any one event) were assumed to be outside a reasonable range and were therefore omitted. These predictions accounted for 0.3% of all predictions and 5.9% of the predictions greater than 0.5 flood report made in the evaluation phase. On the other hand, since the Random Forest model predictions are the average of each regression tree's prediction, the Random Forest predictions cannot exceed the range of training values. Therefore, none of the Random Forest predictions were omitted.

Figure 3.3 shows the predictions made by the Poisson regression model in the training and evaluation phases. The predictions made using the training data as input generally follow the one-to-one line, while the predictions made using the unseen, evaluation data as input are much more scattered. Additionally, for many of the values of true floods, the predicted number of floods in the evaluation phase had large standard deviations (mean of 18.42 flood reports) compared to the training phase (mean of 4.99 flood reports). For some values of true flood reports in the evaluation phase, the range of predictions was large even when the mean of the predictions was close to the true value. For example, the mean of the predictions when there were 31 true flood reports was 29, however, the predictions ranged from 4 to 91.

One explanation for the limited performance of the Poisson regression may be due to the target data not conforming to the assumptions used in the development of the Poisson regression model. One of the primary assumptions when using Poisson regression is that the variance and the mean of the output dataset are equal. In the flood reports dataset described in Section 2.2 and Figure 3.2, the mean was 1.2 flood reports, much lower than the variance, 108 flood reports, meaning that the data were overdispersed. A common

Table 3.3: Summary of training and evaluation results for Poisson regression and Random Forest. All units are in number of flood reports.

		RMSE		MAE		std	
		Training	Evaluation	Training	Evaluation	Training	Evaluation
All days	Poisson	2.31	6.71	0.46	0.96	4.99	18.42
	RF	1.86	3.87	0.30	0.69	3.06	6.00
Non-zero flood days	Poisson	10.06	29.81	6.56	16.34	5.18	19.11
	RF	8.04	16.55	4.41	9.83	3.17	6.21

method for handling overdispersed data in such cases is to use a modified version of Poisson regression called overdispersed Poisson. With overdispersed Poisson the assumption that the mean is equal to the variance is relaxed (Gardner et al., 1995). The overdispersed Poisson was tested as well but the results were very similar to the Poisson regression results.

From the sensitivity analysis of the Random Forest parameters, the number of trees per forest was 100. As seen in Figure 3.4 Random Forests with more than 100 regression trees saw minimal improvements in terms of RMSE, MAE, and std. The model was more sensitive to the number of variables per tree. The number of variables per regression tree that performed the best in the optimization procedure was 16. Changing the number of variables per tree from six (the default) to 16 decreased the models RMSE by 23%.

The RMSE, MAE, and std of training and evaluation predictions from the Random Forest model are reported in Table 3.3. The RMSE was significantly higher in the evaluation phase compared to the training phase both when considering all of the events and when considering only events where floods were recorded. In both cases, the evaluation RMSE was about two times the training RMSE, suggesting that, like the Poisson regression, the model was overfit to the training data. Figure 3.5 shows the predicted number of flood reports made by the Random Forest model in the training and evaluation phases.

One reason for the overfitting seen in the models may be related to the imbalance of dataset. As would be expected, there are far more storm events where no flooding is reported, making the dataset imbalanced. The ratio of storm events for which some rain fell and zero flood reports were made to storm events on which some rain fell and at least one flood report was made is approximately 18:1. Another factor may be the relatively small nature of the dataset (less than 1,000 total records). He and Garcia (2009) noted that models trained with datasets that are both imbalanced and small are particularly prone to overfitting to the training data. As more data is collected and available for use in model training, it is expected that model overfitting would decrease.

Comparing the Poisson regression and Random Forest model, Random Forest performed better overall. In terms of RMSE and MAE, both models were nearly equal in the training phase, however, Random Forest performed significantly better in the evaluation phase. While both models showed signs of being overfit to the training data (i.e. a large drop in performance during the evaluation phase), the proportional difference in performance between training and evaluation in the Random Forest predictions was roughly two-thirds of that of the Poisson regression. A more significant difference in performance between Poisson regression and Random Forest was seen in the stability of the predictions in the evaluation phase as measured by the standard deviations of the predictions. Quantitatively, in the evaluation phase, the standard deviations of the Poisson predictions was more than three times that of the standard deviation of the Random Forest predictions. Visually, the difference is apparent when comparing the standard deviation bars in the evaluation predictions in Figures 3.3 and 3.5.

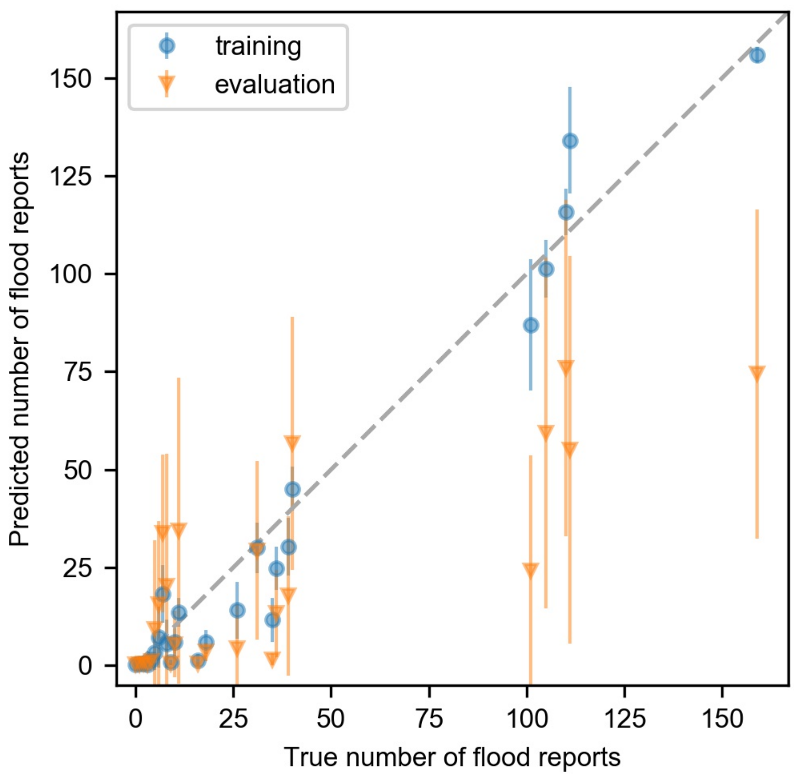

Figure 3.3: Model results for Poisson regression. Error lines represent the standard deviation of predictions.

It is important to note that model performance was measured using observed environmental conditions as model input. In practice, forecasted environmental conditions would be used as model input to predict flood severity. For example, rather than using rainfall and tide data observed at monitoring stations as input, rainfall forecasts from the High-Resolution Rapid Refresh (HRRR) model (Smith et al., 2008) and NOAA's tide predictions (NOAA, 2017d) could be used as inputs. The use of uncertain forecast rainfall data are expected to increase the overall uncertainty of the model (Collier, 2007; Bartholmes and Todini, 2005).

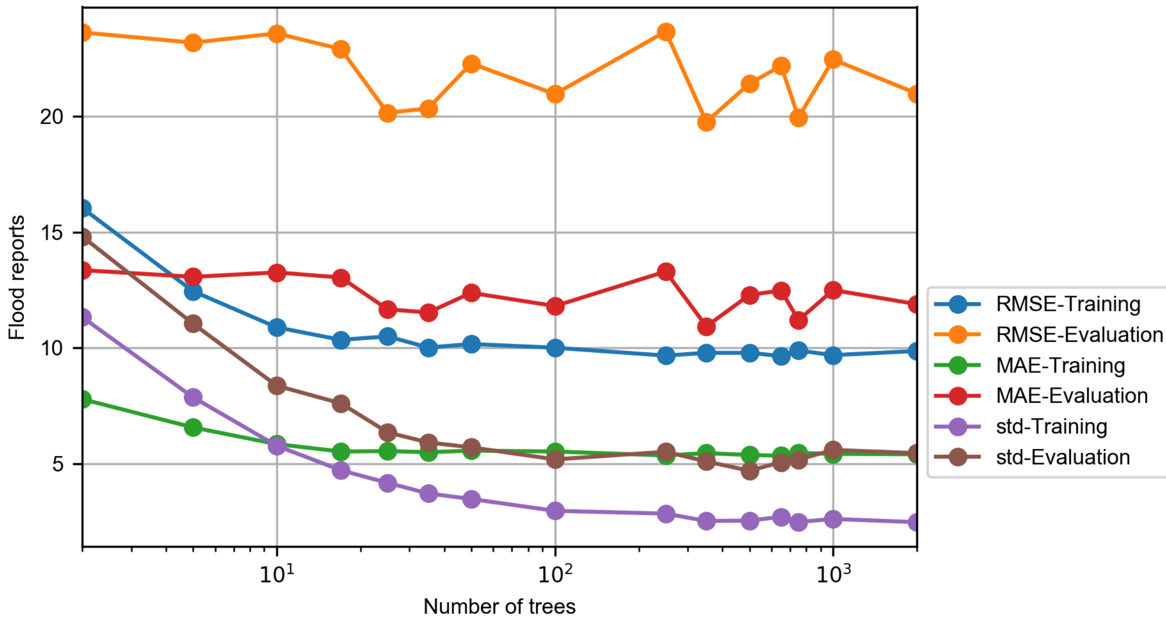

Figure 3.4: Random Forest model results with varying numbers of trees per model.

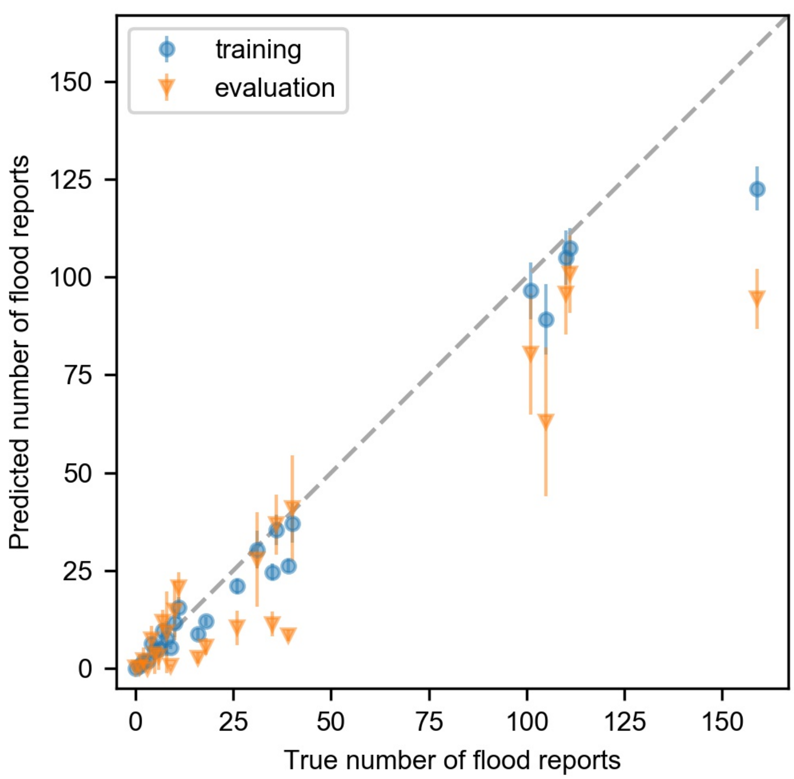

Figure 3.5: Model results for Random Forest. Error lines represent the standard deviation of predictions.

Table 3.4: False positive and false negative statistics for Poisson regression and Random Forest. The statistics in the false positive columns describe the predicted flood reports greater than 0.5 when the true number of flood reports was zero. The statistics in the false negative columns describe the true non-zero flood reports when the predicted flood reports were zero.

	False Positives		False Negatives	
	Poisson	RF	Poisson	RF
rate (%)	4.39	7.09	45.92	34.46
count	1023.00	1653.00	524.00	428.00
25%	0.69	0.70	1.00	1.00
50%	1.10	1.17	2.00	2.00
75%	2.40	2.81	5.00	3.00
max	122.48	36.37	101.00	9.00
mean	3.41	2.47	5.77	2.64
std	8.55	3.18	11.55	2.30

3.4.2 False Negative and False Positive Predictions

Statistics summarizing false negative and false positive predictions in the evaluation phase are given in Table 3.4. Poisson regression had fewer but more variable and extreme false positive predictions compared to Random Forest. The false positive rate for Poisson regression was 4.39% compared to 7.09% for Random Forest, however, the standard deviation of the false positive predictions was much larger from the Poisson regression (8.55 flood reports compared to 3.18 flood reports). For both models, most of the false positive predictions were less than 1.17 flood reports, which would round down to one. The maximum false positive prediction was much greater in the Poisson regression compared to the Random Forest (122 compared to 36).

Compared to the false positive rates, both Poisson regression and Random Forest had much higher false negative rates (45.92% and 34.46% respectively). Importantly, compared to Random Forest, Poisson regression predicted false negatives when true flood reports were higher on average (mean of 6 flood reports compared to 3 flood reports). As with the false positives, the standard deviation of the false negatives predictions was higher for Poisson regression versus Random Forest, 11.55 flood reports compared to 2.30 flood reports. Similarly, the maximum true number of flood reports when a false negative prediction was made was much higher from Poisson regression (101 flood reports compared to 9 flood reports).

Variable Importance

Figure 3.6 shows the importance of each of the input variables as calculated from the Random Forest model in terms of the percent increase in mean squared error (MSE) when each of the variables is permuted individually. The total cumulative rainfall value was by far the most important of the variables. This was much more important than any of the other variables derived from the raw rainfall data, including the maximum hourly and maximum 15-minute rainfall values, suggesting that, in this record, large rainfall volumes caused more flooding than high rainfall intensities. The next three variables in terms of prediction importance were related to tide: low tide, lower low tide, and higher high tide.

The variable importance results shown in Figure 3.6 are supported by the raw data shown in Figure 3.7. The number of flood reports has clear positive relationship with the total cumulative rainfall. The same is true for low tide and lower low tide. The relationship is less clear for the maximum hourly rainfall and the maximum 15-minute rainfall, but according to Figure 3.6, the Random Forest model was still able to glean some meaningful information from these variables. Interestingly, the tide level during the maximum 15-minute rainfall, has a clearer visual relationship with the number of flood reports compared to the maximum hourly, and maximum 15-minute rainfall values, but is considered less important by the Random Forest algorithm. One explanation for this is that the information provided by the tide level during the maximum 15-minute rainfall is already provided to the model, perhaps in a more useful form, from the low tide, lower low tide, and higher high tide variables.

The average height of the water table during a given event, surprisingly, did not add appreciable predictive power to the model. This may suggest that the surficial groundwater table did not impact flood severity in a significant way. However, the fact that it did not provide predictive power does not necessarily mean that the surficial groundwater table level did not contribute to flooding. For example, since the infiltration of rainfall causes the surficial groundwater table to rise, it is possible that the information provided by the rainfall data provides similar but clearer predictive power to the model compared to the surficial groundwater table level.

Although total cumulative rainfall is clearly the dominant predictor of flood severity in this dataset, it is commonly understood that other factors can have a significant impact on flooding in a coastal environment. For example, high tides alone can cause flooding in coastal cites (Marfai et al., 2008). When tidal information is omitted from the model inputs in this study, the RMSE of the Random Forest predictions increases by 4% and by 14% for the Poisson regression predictions. Thus, while rainfall is clearly the most important variable, tide levels and potentially other environmental variables cannot be ignored. It is anticipated that as sea levels rise, the importance of tide levels and water table level in predicting flooding will grow (Hoover et al., 2016).

3.4.3 Potential Explanations for Model Limitations

A likely reason for the limited performance in both models is the limited amount of reported street floods used to train the models. The crowd-sourced flood report dataset used in this study is a unique and valuable dataset, but still a complete picture of flooding impacts is missing. Flood reports were made on only 42, or just over 5%, of the records used in the modeling. In addition, the number of flooding reports were distributed very unequally among the 42 storm events on which flood reports were made. More than 65% of the total

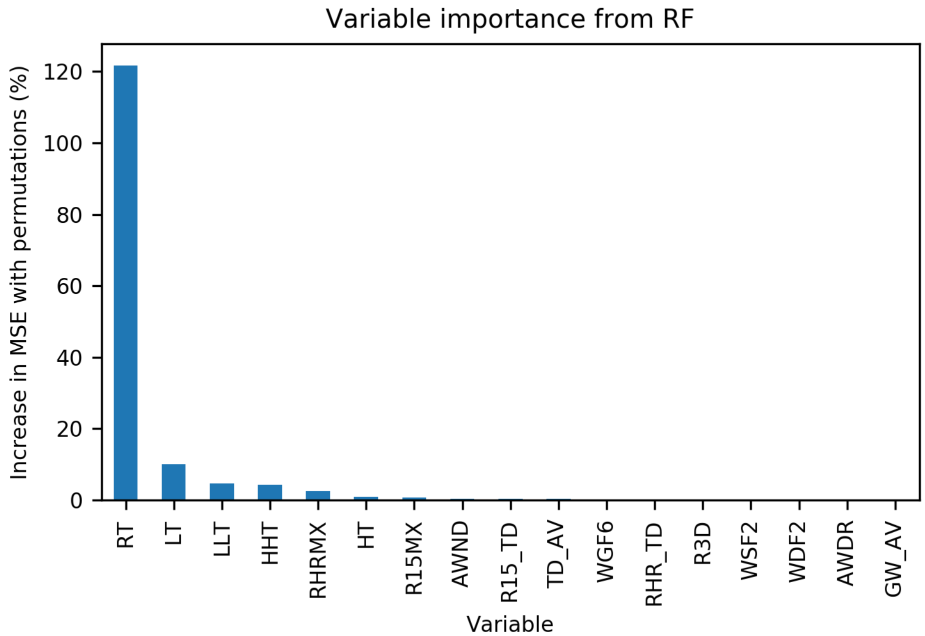

Figure 3.6: Importance of input variables

flood reports were recorded from just six storm events (0.6% of the total storm events modeled). The rarity of storm events with any flood reports, and especially a large number of flood reports, makes it difficult for the model training. Solomatine and Xue (2004) faced similar problems in training their machine learning model to accurately predict high peak flows which occurred rarely in their dataset.

The results also suggest that, compared to storm events with large volumes of rainfall which caused flooding, other types of storm events were not as well modeled by the data-driven models. Figure 3.8 shows the percent error of the Random Forest evaluation predictions for the 11 events with the top 10 number of reported floods (two events had 101 flood reports: Hurricane Nicole and an unnamed event occurring on 20 September 2016). Two unnamed events, heavy rain which occurred on 16 May 2014 (35 reported flood locations) and thunderstorms which occurred on 10 July 2014 (39 reported flood locations) have average percent error magnitudes larger than the rest. Both of these events had much lower cumulative rainfall and tide levels, the most important variables in the model (see Figure 3.6), but higher maximum hourly rainfall and relatively high maximum 15-minute rainfall values compared to the other high flooding events. Given this, it is possible that these events caused flash floods. The worse performance of the models at predicting the flooding severity from these two events may suggest that this type of flooding is not well represented in the training dataset. It is expected that the data-driven models would better predict such flood events with a larger, more complete dataset, containing more instances of similar flooding events. Additionally, with more training data, the model could be trained on specific subsets of flood events tailoring it to a type of flood

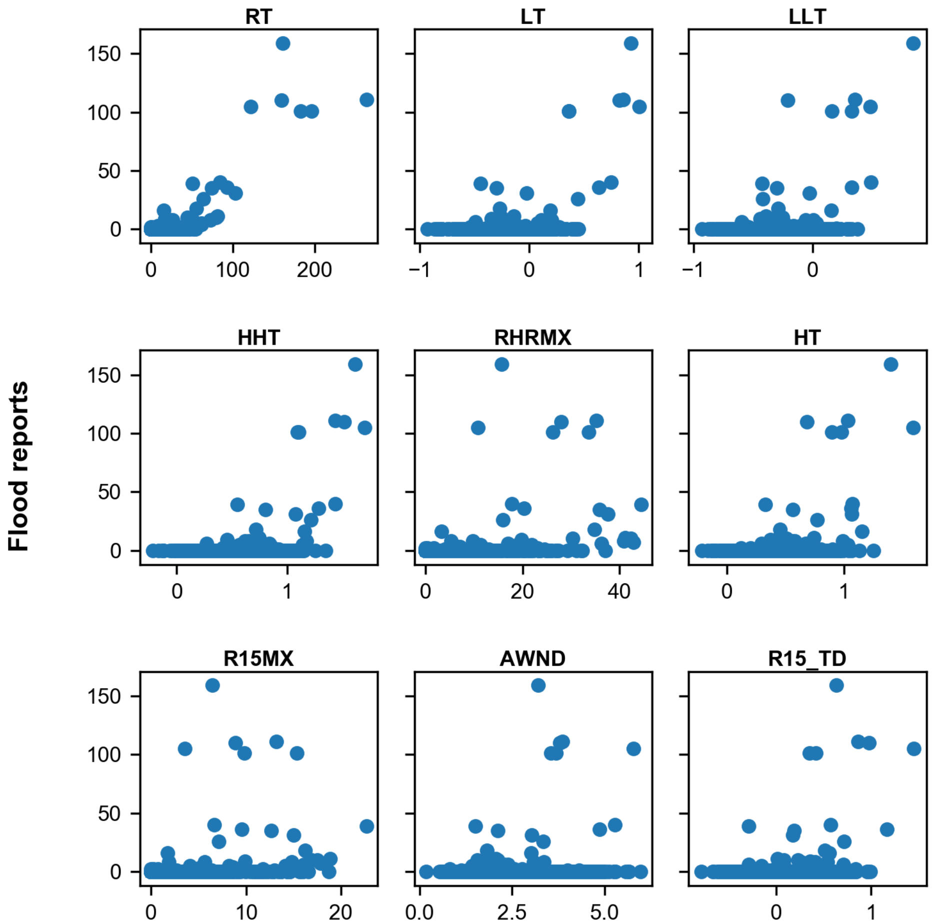

Figure 3.7: Flood reports against top nine predictor variables. Units for each variable are shown in Table 3.2

event with specific characteristics (e.g., flash floods).

Besides the limited number of flooding events with which to train the models, bias present in the training data could have hampered model performance. Because the flooded locations were reported by individuals, there is an unknown amount of subjectivity and bias in the data as can be expected when using crowd-sourced data. Since the models are trained on data reported by individuals, one individual may influence the trained model disproportionally. Figure 3.9 shows the total number of flood reports made by each individual reporter, and the number of flood events for which each reporter recorded at least one flooded location during the period of record. The highest number of total reports made by one reporter was 158, 14% of the sum total reports from all 71 reporters. Therefore, the models in their training, are significantly influenced by this one reporter and can inherit, to some extent, his/her biases.

Another potential bias is in the under- or over-representation of different roadway types in the flooding record. Figure 3.10 shows the percentage of roadway length per VDOT roadway class in Norfolk and the

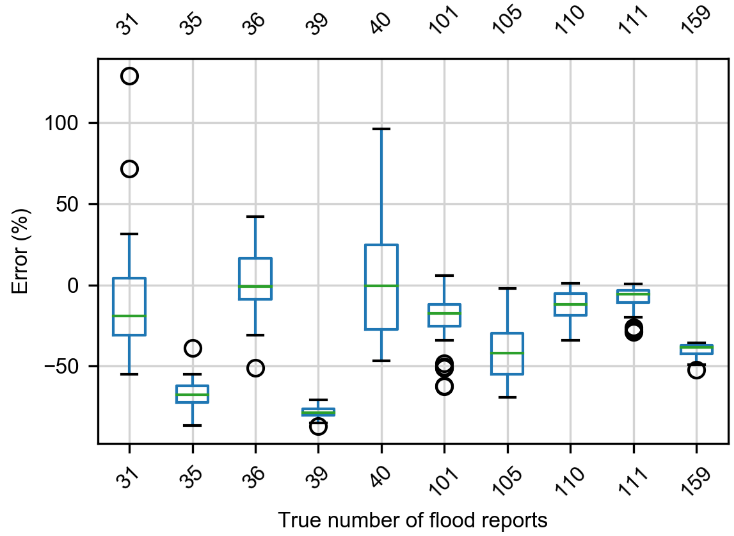

Figure 3.8: Top 10 flooding event percent error from Random Forest evaluation results

percent of each roadway class at which flood reports were made (Table 3.5 gives the descriptions for each of the classes). From the figure, it is seen that although public local streets (class 6) account for the majority of the roadway length of the city (close to 60%), only 40% of the flooded streets reported were public local streets. Conversely, principal arterials (class 3) accounted for nearly 30% of the flooded street reports even though these streets make up less than 10% of Norfolk's total roadway length. This suggests that a flooded street less traveled and, therefore, less important to the overall connectivity of the city's street network, may have flooded but may not have been reported within the record with the same frequency as the more major roads.

A third example of bias may occur when unequal attention in reporting is given to certain geographic areas of the city or to certain storm events. One example of this bias is seen in the difference in reported floods between Hurricane Hermine and Hurricane Matthew which occurred only one month apart. For Hurricane Hermine, 22 flood reports, more than half of the total of 40 flood reports made as a result of Hurricane Hermine, came from one area in downtown Norfolk. In contrast, for Hurricane Matthew, which produced more than three times as much total cumulative rainfall on average than Hurricane Hermine (264 mm compared to 84 mm) and was at least comparable in terms of tide, water table height, and wind conditions, only six flood reports were made from the same area. It is unlikely that the actual flooding caused by Hurricane Matthew, a much larger storm, was in fact a quarter in severity, but more likely that there were significant differences in reporting between the two events.

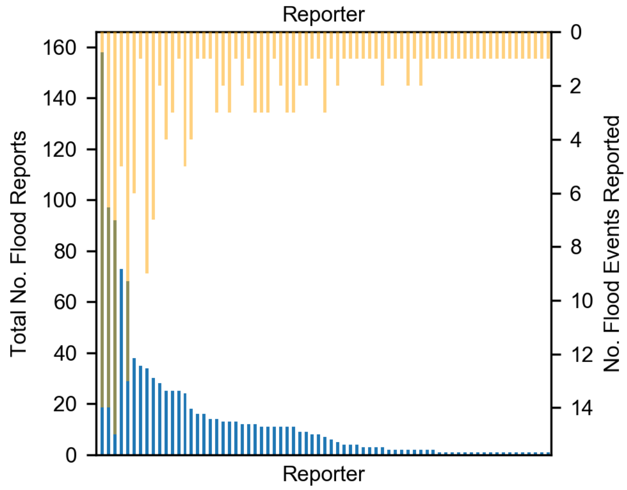

Figure 3.9: Number of total reports made and events reported per reporter

Table 3.5: VDOT roadway class codes and descriptions

VDOT Road Class Code	Description
1	Interstate
2	Tunnel Roads and other VDOT owned
3	Principal Arterials
4	Minor Arterials
5	Collectors
6	Local Streets- Public
7	Local Streets- Private
8	Miscellaneous
9	Base Roads
10	Public Alleys

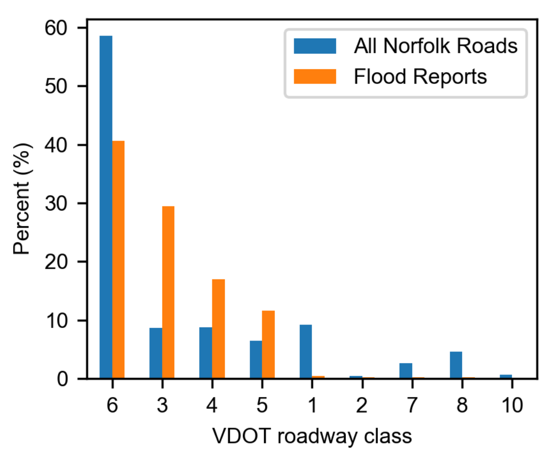

Figure 3.10: Percentage of total roadway length and percentage of reported floods per VDOT roadway class in Norfolk, VA

3.4.4 Increase in Street Flood Reports

Flooding reports and events have become more frequent over the period of record (September 2010 to October 2016). The number of flooded street reports has increased year to year in the past four years and overall in the past seven years (see Figure 3.11). More than twice as many floods were reported in 2016 compared to 2014. Very few flood reports were made in 2013 compared to the other years of record. This can be explained, at least in part, because 2013 was an exceptionally mild hurricane season, the first since 1994 without any major hurricanes. The only storm event in 2013 reported to have caused flooding was an unnamed heavy rain event.

The overall increase in the number of flood reports over the period of record was due primarily to an increase in the number of storm events from which flood reports were made rather than an increase in the number of reports per storm event. In the years 2010-2013, five total storm events were reported to have caused flooding, while in 2014 alone, 16 events were reported to have caused flooding and at least 10 storm events per year resulted in reports of flooded streets in 2015 and 2016 (see Figure 3.11). In contrast to the storm events reported in 2010-2013, most of the storm events reported to have caused flooding in the years 2014-2016 were smaller, unnamed storm events. In each of the years 2010, 2011, and 2012, there was a storm event for which more than 100 flood reports were made, each of them named, major hurricanes (Nicole, Irene, and Sandy, respectively). In 2014, on the other hand, of the 16 storm events reported to have caused flooding, the maximum number of flood reports for an individual storm event was 39, and only one named hurricane was reported to have caused flooding, Hurricane Arthur.

The increase in flood reports due to smaller storms from 2014-2016 may suggest that the City of Norfolk is becoming more susceptible to flooding from less extreme storm events (i.e. storm events which are not hurricanes). It is noted that most of the flood reports were made by staff of the City which operates with limited resources. In reality, the number of flood locations resulting from the recorded storm events could be larger than what was reported, including street floods that may have occurred from less extreme events in 2010-2013. It is possible that the increase in flood reports over the period of record may simply be due to an increase in attention given to street flooding and resources to street flood reporting in Norfolk, rather than an increase in actual flooding. An increase in attention and resources allocated to street flooding and street flood reporting from the City of Norfolk however, may still suggest that flooding problems are worsening.

Sweet and Park (2014) noted an increase of nuisance level tidal flooding in Norfolk, VA from 1.2 days per year in the years 1956-1960 to 7.4 days per year in the years 2006-2010. Sweet and Park (2014) also predicted that the amount of flooding will increase with time due to predicted sea level rise. This is in agreement with the increase of flood reports seen in the present study. As sea level rises and climate changes, it may be necessary to incorporate a mechanism to account for this change into the data-driven models. This change be referred to as concept drift (Gama et al., 2014; Gama and Castillo, 2006) and mechanisms for accounting for it would be especially useful when using Random Forest, which as noted above, cannot exceed the range of training data.

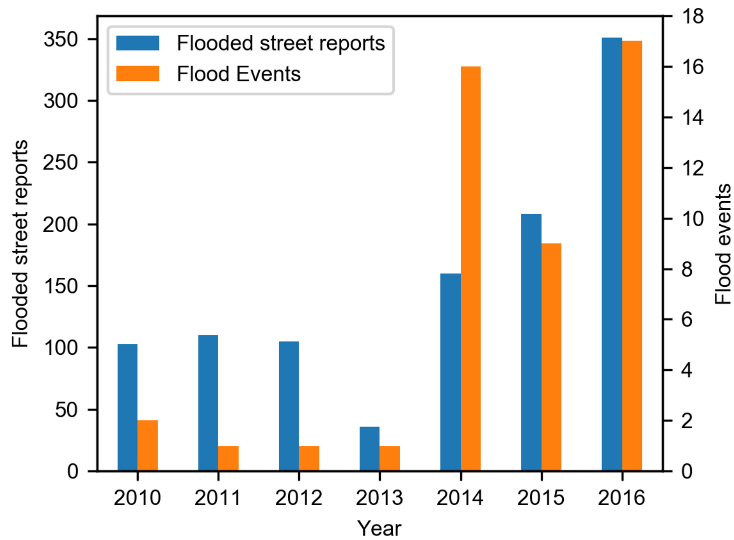

Figure 3.11: Flood events in Norfolk, VA Sep. 2010 - Oct. 2016

3.5 Conclusions

Two data-driven models, Poisson regression and Random Forest were trained to predict flood severity for a given set of environmental conditions (rainfall, tide levels, groundwater levels, and wind conditions) using quality-controlled, crowd-sourced street flooding reports as a proxy output variable. The data used for training and evaluating the models was from Norfolk, Virginia USA. The Random Forest model performed better overall compared to Poisson regression in the evaluation phase (root mean squared error of 3.87 compared to 6.71 flood reports, mean absolute error of 0.69 compared to 0.96 flood reports) with less variance (standard deviation of 6.00 compared to 18.42 flood reports). The most important variable in predicting model output in the Random Forest model was by far total cumulative rainfall followed by low tide and lower low tide.

The quality-controlled crowd-sourced record provided by the City of Norfolk, despite limited coverage spatially and from storm to storm, provided an uncommonly detailed flood record. In the record, flooding at individual intersections and streets was recorded for many events over an extended period of time. Using this as training data, the models demonstrated in this paper could give city workers a reasonable estimate of flooding severity based on forecasted environmental conditions. This is a first step in the long-term goal of spatially and temporally detailed urban flood predictions to assist city managers in real-time flood adaptation measures such as traffic management. This also demonstrates one way that crowd-sourced data, despite its limitations, can provide useful information to flood prediction models.

A main limiting factor in building accurate models is the quantity and quality of the record of flooding used to train the model. Given the bias present in the training dataset, predictions were necessarily lumped spatially (predictions were made at the city scale) and temporally (predictions were made at a event time scale). While other work has raised the need for accurate and dense measurements of rainfall (Sadler et al., 2017; Hill et al., 2014), the primary input to flood models, the results of this work highlight the need for more accurate and complete record of flooding data including depth and duration of flood occurrences. Such data is needed to adequately train flood models with enough spatial and temporal detail to help make street-level, real-time operational decisions.

Given more complete and objective flood occurrence data, it is likely that a data-driven model such as the ones demonstrated in this paper, could predict street flooding with much greater precision. The need for a more complete flood record data may be filled with a street-level sensor network, eliminating human subjectivity. Such a network is currently being piloted to record water levels at commonly flooded intersections in Norfolk. The detailed data from this network could be used to further improve predictions from models such as those demonstrated in this paper. Crowd-sourced data such as flood reports made from cellular devices could also be useful. Although the subjectivity in publicly crowd-sourced data would likely be similar to the dataset used in the paper, a wider number of reporters would presumably mitigate the subjectivity to some degree.

3.6 Data Availability

Data used for the analysis in this study can be found on HydroShare: https://www.hydroshare.org/resource/9db60cf6c8394a0fa24777c8b9363a9b/.

3.7 Acknowledgments

I gratefully acknowledge the Mid-Atlantic Transportation Sustainability Center University Transportation Center for funding of this work. I also thank the Hampton Roads Sanitation District for access to their data. Lastly I thank the City of Norfolk for providing the reported flood data.

At the time of publishing this dissertation, this chapter was in revision with *Environmental Modelling and Software*.

Chapter 4

Leveraging Open Source Software and Parallel Computing for Model Predictive Control of Urban Drainage Systems using EPA-SWMM5

4.1 Introduction

Researchers have predicted that storm intensity will increase on average due to climate change (Berggren et al., 2012; Neumann et al., 2015). Coastal cities have an additional challenge as sea levels rise which makes it more difficult to drain storm runoff from streets. Coastal cities have already experienced increased flooding from high tidal events alone (Sweet and Park, 2014).

More intense storms and rising sea levels will put greater stress on urban drainage systems necessitating changes for urban drainage systems to perform at current levels. One possible adjustment is to make capital improvements such as increasing pipe size or constructing new storage units. Another option is to convert drainage systems from passive, gravity driven systems to active or "smart" systems (Kerkez et al., 2016). Active systems can increase performance of a urban drainage system at a lower cost than traditional capital improvements (Meneses et al., 2018). Actively controlling an urban drainage system does not increase the *actual* capacity of urban drainage infrastructure, but rather more efficiently uses the existing infrastructure, increasing its *effective* capacity. For example, one part of an active urban drainage system could be a valve at the outlet of a retention basin which can be automatically opened or closed based on system conditions and forecasts. With this setup, the valve could be closed more during a storm which would utilize the available storage better than would have been possible without the valve.

For an active urban drainage system to achieve its objective (e.g., minimize flooding, reduce combined-sewer overflows), an effective management strategy is required. Management decisions for a urban drainage system include which actuators (e.g., valves and pumps) in the system should change, when to change them, and to what setting. I refer to these decisions as a control policy (Vrabie et al., 2009; Mayne et al., 2005; Langson et al., 2004). An effective control policy for an active urban drainage system may depend on a number of factors such as antecedent moisture conditions, expected intensity and duration of oncoming rainfall, current water levels in the system, the condition of the drainage infrastructure, and other factors (e.g., tide levels in tidally influenced urban drainage systems).

A common approach for determining an effective control policy is model predictive control (MPC) (Camacho and Bordons, 2007). MPC has been used effectively in many control applications including automotive controls (Del Re et al., 2010), HVAC (heating, ventilation, and air conditioning) (Afram and Janabi-Sharifi, 2014), and other industrial applications (Qin and Badgwell, 2003). MPC has also been used effectively in urban drainage applications (Puig et al., 2009; Cembrano et al., 2004; Schütze et al., 2004; Gelormino and Ricker, 1994). In MPC, a process model is used to simulate the physical system and evaluate alternative control policies. Forecast data can be used as input for the simulation. During the control period, on-line optimization is performed, meaning that an optimal control policy is found and implemented at each time step (Camacho and Bordons, 2007).

Although effective for finding effective control policies, implementing MPC for a urban drainage system is non-trivial due to the dynamics within the system. The fundamental governing equations for modeling urban drainage systems are the St. Venant equations which, when considered fully, are non-linear (Tayfur et al., 1993). This makes finding an optimal control policy for urban drainage systems challenging using MPC (Darsono and Labadie, 2007). To address this dilemma, two alternative approaches have been used. The first is to simplify the governing equations of the process model to a linear system. This makes the optimization problem solvable using well-established procedures such as simplex (Nelder and Mead, 1965). Gelormino and Ricker (1994) took the approach of linearizing their system, converting their process model into a linear-time-invariant model to perform MPC for a large combined sewer system in Seattle, Washington USA.

The second approach is to retain the non-linear St. Venant equations and use a metaheuristic to find the best control policy at each time step. In this approach, a true optimization procedure is not possible because the system remains non-linear; instead, a metaheuristic (e.g., an evolutionary algorithm (EA)) can be used

(Gandomi et al., 2013). The use of a metaheuristic precludes the possibility of determining a guaranteed optimal control policy and is typically computationally expensive. The advantage of this approach, however, is that the non-linear governing equations in the process model are retained. This approach was taken by Heusch and Ostrowski (2011) who used a dynamically dimensioned search for finding the best control policy and the United States Environmental Protection Agency's Stormwater Management Model Version 5 (EPA-SWMM5), which numerically solves the St. Venant equations, as their process model (Huber et al., 2005). Similar to Heusch and Ostrowski (2011), I have selected to follow the second approach so that the non-linearities in the process model can remain, and to leverage EPA-SWMM5 as the process model.

EPA-SWMM5 is an attractive choice as a process model for urban drainage systems for several reasons. EPA-SWMM5 is in the public domain making it free of charge and its source code is open-source making it customizable. The model simulates a wide variety of urban drainage structures including active controls such as orifices with variable openings and pumps. EPA-SWMM5 has been used in many research applications, as well as in engineering practice to model urban drainage systems (Burger et al., 2014). Notwithstanding the wide use and utility of EPA-SWMM5 for modeling urban drainage systems, and the established utility of MPC as a successful approach for determining effective control policies there is currently no software available for performing MPC using EPA-SWMM5. Although Heusch and Ostrowski (2011) developed software that implements MPC with EPA-SWMM5, that software was closed-source and is no longer available.

This study advances the work done by Heusch and Ostrowski (2011) by creating an open-source implementation of MPC for EPA-SWMM5, *swmm_mpc*, and by demonstrating swmm_mpc's parallel computing capabilities. By making swmm_mpc open source, other researchers will be able to use, improve, and build from the source code. Although, the software written by Heusch and Ostrowski (2011) supported the use of parallel computing, this capability, which is critical to the usability of such software given its associated computational costs, was never demonstrated or tested in the literature.

swmm_mpc was written in the Python programming language. Several third-party Python packages were necessary for the success of this project including pyswmm (https://github.com/OpenWaterAnalytics/pyswmm) and the Distributed Evolutionary Algorithms for Python (DEAP) (https://github.com/DEAP/deap). To evaluate swmm_mpc, it was applied to a use case model with two active control devices. The swmm_mpc results were compared to the results from a rules-based approach and a scenario with no active control. The swmm_mpc software was run on a desktop personal computer (PC), a high-performance computer (HPC), and a rented, cloud-based machine to demonstrate and test the parallel processing capability of the software.

The remainder of this paper describes the methods used to implement swmm_mpc including a description of the MPC workflow and the interaction and role of the third-party Python libraries. The use case model is then described and the results of the evaluation are presented and discussed. As part of the results and discussion, the benefits of parallelization and the use of a high-performance and cloud-based computing for running swmm_mpc are quantified and discussed.

4.2 Methods

4.2.1 Overview of MPC for urban drainage systems

MPC for a urban drainage system consists of three main components as shown in Figure 4.1. The first component is the physical system, including the system states and system controls. The system states include hydraulic states such as water levels at system nodes and flow rates in system pipes, and hydrologic states such as watershed soil moisture and runoff. In a real system, these states would come from real-time sensors. The system controls are actuators that accept and implement the settings resulting from the MPC process at each time step.

The second component in MPC is a process model used to simulate the future states of the urban drainage system. The process model uses the states read from the urban drainage system as its initial states. The process model also takes future model inputs such as rainfall or tide level. Given the current state of the system and future disturbances, the process model is used to evaluate the effectiveness of control policy candidates.

A control policy consists of one setting for each actuator, for each control time step, for the duration of the control horizon. An individual setting can be a number as would be the case for a valve where the number would correspond to the percent open of the valve. An individual setting can also be a binary setting as would be the case for a pump that can either be on or off. As an example, consider a system with a variably opening valve and an on-off pump with a control horizon of three hours and a control time step of 15 minutes. A control policy for this system would consist of two arrays, an array of numbers between 0 and 1 to specify the percent open of the valve should be, and an array of "on" or "off" to specify the setting of the pump. Both arrays would have 12 settings (four settings per hour for three hours).

To evaluate the effectiveness of a given control policy, the settings in the policy are applied to simulated actuators in the process model and the process model simulation is executed. At the end of the simulation, a cost is determined for the policy. The cost is based on a user-defined cost function. In this study, I consider mainly the cost resulting from flooding but other costs could be considered within this general framework including the costs of combined sewer overflows (CSO) and water quality. The cost may also be a factor of other process model outputs such as deviation from target water levels at certain points (Schütze et al., 2004).

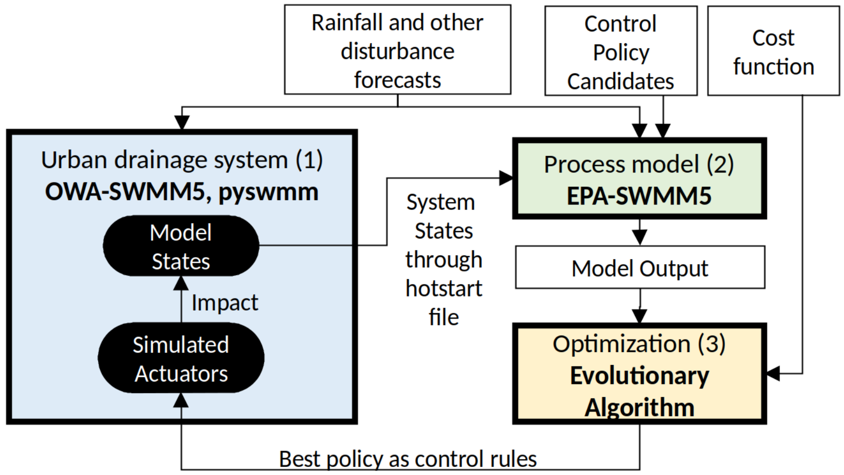

Figure 4.1: Main components of MPC in swmm_mpc system

The third component of MPC for a urban drainage system is an optimization routine to determine the best control policy for the system. Using the process model to assign a cost to a given control policy, the optimization procedure seeks to find the control policy that incurs the smallest cost. If the process model is linear, a true optimum can be found using traditional optimization procedures like simplex (Nelder and Mead, 1965). If the process model is non-linear, other approaches must be taken such as using a metaheuristic to find an effective control policy (Gandomi et al., 2013).

In summary, the chronological workflow for MPC for a urban drainage system is: 1) system states are read from the physical system, 2) using the system states as initial conditions and future disturbances as input, a process model is used to evaluate control policies, 3) the best control policy is selected through an optimization procedure, and 4) the best control policy is implemented in the real system. Although the best control policy is obtained for the entire control horizon, only the first step in the control policy is used since the procedure re-optimizes at every control time step.

4.2.2 MPC for SWMM5: swmm_mpc

Implementation of the parts of MPC using Python and SWMM5 was done in the swmm_mpc Python package and uses the standard EPA-SWMM5 and an enhanced version of SWMM5 developed by Open Water Analytics, OWA-SWMM5. The software simulates online MPC for an urban drainage system using SWMM5 as the process model and as the simulated physical system. The current system could also be used in an offline mode where a control policy for a forecast storm event is found beforehand.

Simulated urban drainage system: OWA-SWMM5 and pyswmm

SWMM5 was used to simulate the physical urban drainage system. For this, an enhanced version of SWMM5, OWA-SWMM5 (https://github.com/OpenWaterAnalytics/Stormwater-Management-Model), was used via an accompanying Python library, pyswmm (https://github.com/OpenWaterAnalytics/pyswmm). Both OWA-SWMM5 and pyswmm were developed and are distributed by Open Water Analytics. Compared to EPA-SWMM5, OWA-SWMM5 contains additional C functions that are accessed by pyswmm.

OWA-SWMM5 and pyswmm provide three key functionalities needed to simulate the online optimization procedure required by MPC. First, unlike when a simulation is run via EPA-SWMM5, when using pyswmm, custom Python routines can be executed between each time step of the simulation. This is critical to swmm_mpc because at each time step in the workflow three processes occur: 1) the states from the simulated urban drainage system need to be read and transferred to the process model; 2) the metaheuristic needs to be run; and 3) the best policy found by the metaheuristic needs to be implemented in the simulated urban drainage system. Using pyswmm, Python code can be run to perform each of these processes at each control time step.

Second, pyswmm enables the transfer of system states at each time step from the simulated urban drainage system to the process model. This is accomplished through a hotstart file. A SWMM5 hotstart file contains all of of the hydraulic and hydrologic states of the model at the time in the simulation when the hotstart file is saved. When a hotstart file is read into a simulation, that simulation's initial hydraulic and hydrologic states are the states represented in the hotstart file. This functionality is well-suited to transfer the states of the simulated urban drainage system to the process model in the swmm_mpc workflow.

Using EPA-SWMM5, a hotstart file can be saved only at the end of a simulation. This is a critical limitation because in MPC the system states need to be transferred at every time step. To address this limitation, I added new functionality to OWA-SWMM5 and pyswmm to enable hotstart files to be saved at any point in a SWMM5 simulation executed using pyswmm. This functionality allowed the system states of the simulated urban drainage system to be transferred to the process model at each time step.

Third, through pyswmm the best control policy found by the metaheuristic can be implemented at each time step. This is done using pyswmm to change the settings of the actuators in the model during the simulation. When a simulation is initialized in pyswmm, each object in a SWMM5 model (every node, link, subcatchment, etc.) can be read into a Python object via its element ID as defined in the SWMM5 input file. Each of these Python objects has attributes that can be read (e.g., depth at a node and flow in a link). Actuators in the model read into Python objects also have the "target_setting" attribute that can be written. To implement a control setting for an actuator via pyswmm, its "target_setting" is set to the first setting in the best control policy.

Process model: EPA-SWMM5

In addition to representing a real urban drainage system, SWMM5 was used as the process model. However, in contrast to using OWA-SWMM5 to simulate the physical urban drainage system, the standard EPA-SWMM5 was used as process models. This was necessary because the current version of pyswmm cannot run more than one simulation at a time. This is a functionality needed in swmm_mpc because at each time step during the simulation of the urban drainage system, at least one process model simulation is run in a predictive fashion to evaluate control policy candidates. EPA-SWMM5, unlike pyswmm, supports multiple simulations being executed simultaneously.

Active controls in EPA-SWMM5

EPA-SWMM5 simulates the active control of certain hydraulic structures including pumps, orifices, and gates. Each of these structures has a setting that can be assigned. For example, the setting for an orifice is a decimal number between 0 and 1 which corresponds to the percent open of the orifice (e.g., a 0.5 setting would mean the orifice was 50%). The user can also define an amount of time for a structure to implement a change in setting. This "time to change" parameter in EPA-SWMM5 represents the delay seen in reality for changing an actuator's setting.

Changing controls during an EPA-SWMM5 simulation is done using one or more control rules (see example in Figure 4.2). A control rule is specified in the SWMM5 input file before the simulation begins and consists of four parts. The first two parts of a control rule are the rule name and the condition. In the example, the rule name is "R1". The condition is "IF NODE J1 DEPTH <2", meaning that the program will check if the depth at the node with the ID of "J1" is less than 2 (the units being defined globally in the model input file as feet or meters). In EPA-SWMM5 the condition can be the state at any link or node and can also be related to global simulation states such as the model simulation time. The third part of a control rule defines which structure(s) should change if the specified condition is met. In the example the structure that will change is "ORIFICE O1" . Finally, the fourth part of the rule defines the setting to which the structure should change. In the example, this is "0.6", meaning that if the condition is met, the orifice should be set to 60% open.

In swmm_mpc a control policy is a time series of control settings (one control setting per control time step for the control duration). This is implemented in EPA-SWMM5 as a set of control rules. In the current version of swmm_mpc, only the control of orifices is supported while support for other controls such as pumps can be added in future versions. Since a control policy in MPC is a time series, each control rule's condition is based solely on the model's simulation time in decimal hours. For example, Figure 4.3 shows a control policy of four settings (0.2, 0.4, 0.5, and 0.2) for "ORIFICE R1" at a 15-minute control time step

```
Rule R1
If NODE J1 DEPTH <2
THEN ORIFICE O1 SETTING = 0.6
```

Figure 4.2: Example of a control rule in SWWM5

```
Rule R1
If SIMULATION TIME < 0.25
THEN ORIFICE O1 SETTING = 0.2

Rule R2
If SIMULATION TIME < 0.50
THEN ORIFICE O1 SETTING = 0.4

Rule R3
If SIMULATION TIME < 0.75
THEN ORIFICE O1 SETTING = 0.5

Rule R4
If SIMULATION TIME < 1.0
THEN ORIFICE O1 SETTING = 0.2
```

Figure 4.3: Example implementation of control policy as set of control rules

implemented as control rules. This text would be written to the EPA-SWMM5 process model input file under the "CONTROLS" heading.

Metaheuristic: evolutionary algorithm

Because I used EPA-SWMM5 as a black-box process model, a metaheuristic was used in place of a true optimization procedure to find an effective control policy at each time step in the MPC run. I chose an evolutionary algorithm (EA) for the metaheuristic since it has been shown to be successful in other urban drainage control applications (Zimmer et al., 2015, 2018) and it's inherent propensity for parallelization Maier et al. (2014). An EA begins with an initial population of individuals where, in my case, each individual is a control policy. A fitness score (or conversely a cost) is assigned to each individual in the population and certain individuals are selected to survive into the next generation based on their fitness score. Mechanisms for improving the fitness of the individuals from one generation to the next mimic natural processes including cross-over and mutation (Maier et al., 2014). The process of selection and improvement is repeated from generation to generation until a stopping criteria is met. Common stopping criteria include a user-defined number of generations or an acceptably low rate of improvement from one generation to the next. The use of an EA requires several user-defined parameters including the number of individuals in the initial population, the cross-over rate, the mutation rate, and the stopping criteria.

Since the EA searches for the policy that incurs the minimum cost, the way in which a cost is assigned to each individual control policy is very influential on the EA's effectiveness. In swmm_mpc, the cost of a control policy is determined using the process model and a cost function. First, each individual control policy is implemented in the process model input file as a set of control rules as described above. Once the control policy is implemented, the EPA-SWMM5 model is executed. Elements of the model output resulting from the process model execution become input for the cost function. The cost function used in swmm_mpc is

$$Cost = \alpha(\boldsymbol{a} \cdot \boldsymbol{v}) + \beta(\boldsymbol{b} \cdot \boldsymbol{d}) \tag{4.1}$$

where $\boldsymbol{a}$, $\boldsymbol{v}$, $\boldsymbol{b}$, $\boldsymbol{d}$ are each 1-dimensional vectors, and α and β are scalers. The members of $\boldsymbol{a}$ are user-defined weight values for flooding at any node in the system and the members of $\boldsymbol{v}$ are the magnitude of flooding at each node as calculated by the process model. The members of $\boldsymbol{b}$ are user-defined weights for deviation from user-defined target water levels at each node in the system and the members of $\boldsymbol{d}$ are the average absolute deviations from target water levels again as calculated by the process model. α, and β, are user-defined constants used to scale and give overall weights to flooding costs compared to deviation costs. Typically weights for the components of the cost (or objective) function sum to 1 and can include a scaling factor to account for variables in different units or scales (Kim and de Weck, 2005). In this formulation, α and β include both the weight and the scaling factors for the objectives.

I intentionally made this cost function flexible so that users can customize it to meet their objectives which may vary between use cases. A cost for flooding is obviously important as that is a major concern for many communities and the prevention of which is one of the main purposes for urban drainage systems. I also included a cost from deviations for target water levels because, in certain cases, it is desirable to maintain water levels close to a target depth. For example, it may be important to keep a certain amount of water in a retention pond for aesthetic and/or ecological purposes. Although the cost function is flexible, when implemented in swmm_mpc, the user need only define what is important to the specific application. For example, default for $\boldsymbol{a}$ is a vector of all 1's. When one node is specified, the weight of any unspecified node becomes zero. The default for $\boldsymbol{b}$ is all zeros, since the user has to specify a target depth for a given node.

To execute EAs I used the Distributed Evolutionary Algorithms for Python (DEAP)(https://github.com/DEAP/deap) library. An advantage of EAs is that they can easily be run in parallel since they performs many independent evaluations (Maier et al., 2014). In DEAP parallel processing is supported through integration with the built-in "multiprocessing" Python library.

swmm_mpc workflow

The MPC workflow in swmm_mpc was implemented using three main Python functions (see Figure 4.4). The function in the workflow called by the user is "run_swmm_mpc." This function runs the MPC workflow and calls the two other main functions. "run_swmm_mpc" takes 13 user inputs as shown in Table 4.1. Through these inputs, the user specifies the model input file to that represents the urban drainage system, the control inputs (i.e., which controls to find a policy for, the control time step, and the control horizon), and EA parameters (e.g., number of generations, cost function parameters).

The most complex of the user-supplied arguments are "target_depth_dict" and "node_flood_wgt_dict" (see Snippet 1 for examples). These two arguments define the $\boldsymbol{a}$ and $\boldsymbol{b}$ variables in the cost function. Additionally, the "target_depth_dict" argument is used to determine $\boldsymbol{d}$. These arguments map from Python data structures to the mathematical variables in the cost function. The "target_depth_dict" argument is a dictionary whose keys are node ids and whose values are dictionaries. The inner dictionary has two keys, the target depth of the node and the weight of the cost for deviations from the weight at the node. In Snippet 1, the "target_depth_dict" specifies that the target depths of Nodes St1 and St2 are 4.0 and 3.5, respectively. The weights are also specified: deviation from the target depth at Node St1 will be twice as costly as deviation from Node St2. The "node_flood_wgt_dict" is a simpler dictionary, the keys of which are node ids and the values are weights. In Snippet 1, the "node_flood_weight_dict" specifies that flooding at Node J3 is five times costlier than flooding at Node St1. Note that if one or more node is included in the "target_depth_dict" or the "node_flood_wgt_dict", other nodes are not included in the cost calculation (in terms of the cost function, the corresponding weights in $\boldsymbol{a}$ and $\boldsymbol{b}$ are zero). This is shown in Snippet 1, the weight of deviations from a water level at Node J3 and the weight of flooding at Node St2 would both be zero since they are not included in the dictionaries.

Snippet 4.1: Examples of "target_depth_dict" and "node_flood_wgt_dict"

```
target_depth_dict = {"Node St1": {"target": 4.0, "weight": 2},
                     "Node St2": {"target": 3.5, "weight": 1}}

node_flood_weight_dict = {"Node J3": 1, "Node St1": 0.2}
```

In the "run_swmm_mpc" function, the SWMM5 model simulating the urban drainage system is run step by step via pyswmm. At the beginning of the simulation, the SWMM5 input file representing the urban drainage system is copied. This copy serves as the input file used for the process model. To ensure that the states and simulation periods of process model remain in sync with the simulated urban drainage system, at each time step a hotstart file from the urban drainage system simulation is saved and then used as the initial states for the process model. The process model's simulation start date and time are also updated to match the urban drainage system simulation's current date and time.

Once the process model's simulation start date and time is same as the simulated urban drainage system and the hotstart file of the simulated urban drainage system is saved, the "run_ea" function is called. The "run_ea" function initiates the EA which starts by creating an initial population of individual control policies. In my case, an individual is a 1-dimensional vector, each member of which is a setting for an individual actuator for one control time step. The initial population for the first time step is a group of random individuals. For subsequent time steps, elitism is used where the best policy found in the previous time step is used to seed the initial population of the current time step.

The control policies initiated in the "run_ea" function are input into the third main function, "evaluate". The evaluate function makes a copy of the process model input file and the input hotstart file. To avoid file naming conflicts, a random string is appended to the hotstart and input file names. The control policy is then implemented in the newly created input file by adding corresponding control rules. Once the control policy is implemented in the input file, the simulation is executed with EPA-SWMM5. When the simulation run is completed, the "evaluate" function parses the output file to determine $\boldsymbol{v}$ and $\boldsymbol{d}$ in the cost function. The policy's cost can then be determined since the remaining cost function parameters (α, $\boldsymbol{a}$, β, and $\boldsymbol{b}$) are user-defined. The evaluation of an individual control policy is independent of all others, therefore, the "evaluate" function is the part of the workflow that is parallelized through Python's "multiprocessing" module.

Using the cost that has been assigned to each policy, the "run_ea" function selects the best individuals to retain in the population of policies for the next generation. After the user-defined number of generations are complete, the best policy found by the EA is implemented in the simulated urban drainage system in the "run_swmm_mpc" function. The policy is also saved and used to seed the population of the next time step. Finally, the setting for that time step is recorded so that at the end of the simulation, the best control policy for the entire simulation time is saved.

4.2.3 System demonstration

To demonstrate the utility and functionality of swmm_mpc, three control scenarios were used and compared for two cases in which the demonstration model and the cost function parameters differed slightly. The computational costs of running swmm_mpc were also quanitfied.

Demonstration model and rainfall event

The model used to demonstrate swmm_mpc (see Figure 4.5) has two orifices (R1 and R2) that control the flow out of two storage units in parallel in the system (St1 and St2, respectively). In SWMM5 models, storage units are generic and are used to represent both natural storage facilities such as a pond as well as man-made facilities such as an underground tank or retention pond. The two orifices from the storage units meet and flow through a junction, J3, before leaving the sytem through the outfall. For the example use case I used a

Table 4.1: User inputs for "run_swmm_mpc" function

Parameter name	Data type	Description	Default value
inp_file_path	String	File path to SWMM5 input file (.inp)	N/A
control_horizon	Number	Control horizon in hours	N/A
control_time_step	Number	Control time step in seconds	N/A
control_str_ids	List of strings	IDs of control structures to be adjusted	N/A
work_dir	String	Path to directory where temporary files will be stored	N/A
results_dir	String	Path to directory where results should be stored	N/A
target_depth_dict	Dictionary	IDs of nodes and corresponding target depths and relative penalty weights	Null
node_flood_weight_dict	Dictionary	IDs of nodes and corresponding relative penalty weights	Null
flood_weight	Number	Overall weight of flood penalties	1
dev_weight	Number	Overall weight of deviation penalties	1
ngen	Number	Number of generations that the GA should perform	7
nindividuals	Number	Number of individuals in the inital GA population	100
run_suffix	String	suffix to be appended to results file	

2-year, 12-hour rainfall event for Norfolk, Virginia, a coastal city that experiences frequent flooding (Mitchell et al., 2013). The rainfall event (see Figure 4.6) had 78.2 mm of total rainfall (Bonnin et al., 2018) with an SCS Type II temporal distribution (Mockus, 2012). The model simulation time was 24 hours.

Control scenarios

Scenario 1: Passive In this scenario there is no active control with the outlets 100% open. Two outlets to the storage units were simulated, one at the bottom of the storage units, the other is near the top. Water continously leaves the retention pond through the lower outlet, and when the water level reaches a certain depth, the water flows out of the upper outlet as well to avoid overtopping.

Scenario 2: Rule-based control For this scenario, a simple logical rule controls the orifice openings and therefore the discharge from the storage units. In practice, such rules can be based on experience and

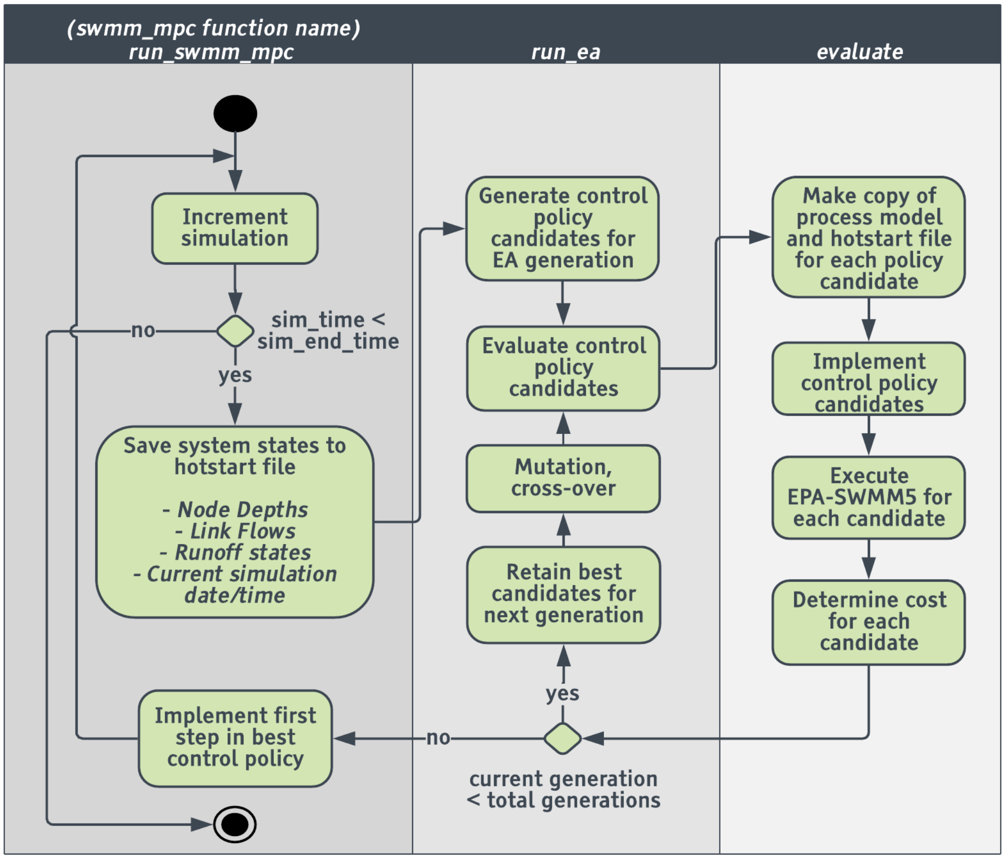

Figure 4.4: Activity diagram of MPC implementation using Python

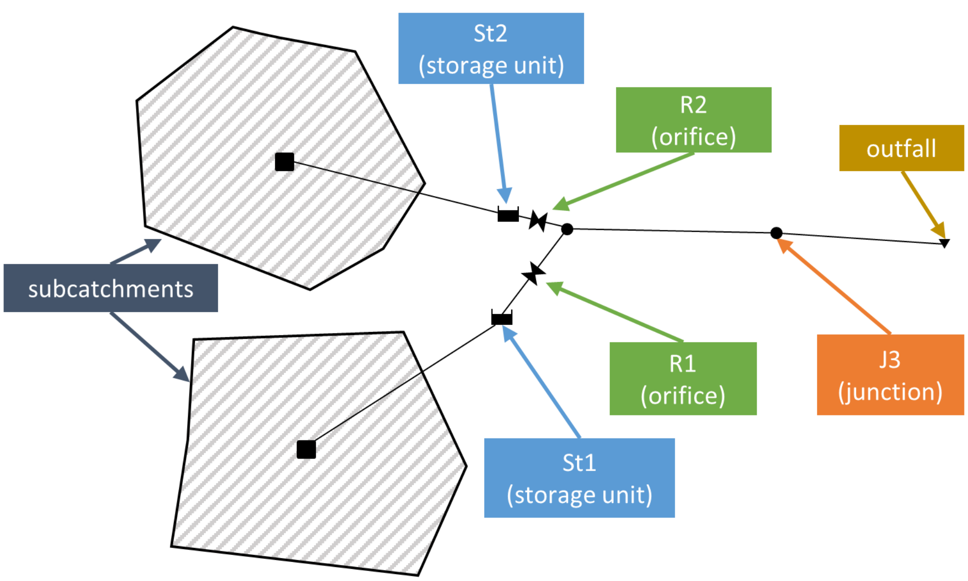

Figure 4.5: Demonstration model schema

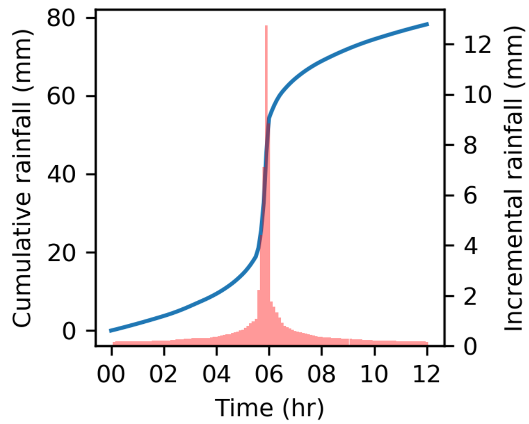

Figure 4.6: Design storm input to use case study

Table 4.2: Control parameters for demonstration cases

Parameter	Value
Control Horizon (hr)	1
Control Time Step (hr)	0.25
Num controls	2

knowledge gained by local stormwater personnel over time. Although, heuristic-based rules alter a dynamic, actuated system, the rules themselves are static, meaning that they do not change for the duration of an event. Furthermore, the rules do not adjust based on forecast conditions. The rule I chose was to set the percent open of the orifices equal to the percent full of the storage units. For example, if one storage unit were 30% full, the valve at the outlet would be set to 30% open. This rule was intended to minimize flooding by retaining some water so as not to flood the downstream nodes, but release enough water to avoid overtopping.

Scenario 3: MPC The MPC control policy was found using swmm_mpc as described in Section 4.2.2 above. One advantage of MPC over the rule-based control is the ability to adjust the actuators based on forecast conditions. For the use cases, I used a control time step of 15 minutes and a control horizon of one hour. Therefore, with two controls, a single control policy consisted of a vector of eight values (2 controls x 4 control steps per hour x 1 hour) (see Table 4.2).

Model and cost function cases

I implemented and compared the three control scenarios for two cases in which model and cost function parameters were slightly different (see Table 4.3). For the passive and rules-based control scenarios, there is only one difference between Case A and Case B. In Case A, the maximum depth of the storage units is 1.52

Table 4.3: Differences in MPC cost function parameters and storage capacity between two test cases

	Case A	Case B
α	3	1
$\boldsymbol{a}$	[0, 0, 1]	[0, 0, 1]
β	0.5	0
$\boldsymbol{b}$	[1, 1, 0]	N/A
Storage depth at St1 and St2 (ft)	1.52 m (5.00 ft)	1.37 m (4.50 ft)

m (5.00 ft) while in Case B the maximum depth is 1.37 m (4.50 ft). The maximum depth of the storage units was reduced in Case B in order to explore the effectiveness of the three scenarios in a more constrained situation.

For the MPC scenario, there is a difference in the cost function parameters between Cases 1 and 2. In Case A, there are two objectives: 1) minimize flooding at the downstream Node J3, and 2) minimize deviations from a target water level of 1.22 m (4.00 ft) at the storage units. These objectives translated to the $\boldsymbol{a}$ and $\boldsymbol{b}$ parameters of the cost function as a vector of zeros except the last member (Node J3), and a vector of zeros except for the first two members (St1 and St2), respectively. In this case, to emphasize minimizing flooding at Node J3 over minimizing deviations from the target water depths at the storage units, I set the cost of flooding weight, α, six times larger than the cost of deviations, β (3 compared to 0.5). In Case B, there was only one objective, minimize flooding at Node J3. In this case, $\boldsymbol{a}$ is the same as in Case A, but β is zero since minimizing deviations from a target water level is not part of the objective.

Use of parallel, high-performance, and cloud computing

The EA used for selecting the best control policy is computationally expensive and therefore, some analysis of computational costs for executing the swmm_mpc workflow was performed. The wall-clock times for Case A (the more complex of the two cases) were compared when using a typical personal computer (PC) and the University of Virginia's high-performance computing (HPC) system, Rivanna. Recognizing that many (likely most) municipalities will not have HPC resources available to them, I also explored the use of a commercial cloud computing service for running swmm_mpc. These services, such as Amazon Web Services, Google Cloud Platform, or Microsoft Azure, allow users to rent large, powerful computers, charging only for the time that the computers are being used. To explore the option of renting a cloud-based machine, I also executed Case A through Google Cloud Platform (GCP). The number of cores available, RAM, and processor speeds of the PC, HPC, and GCP machines are listed in Table 4.4. Case A was run with varying number of cores on each platform.

Table 4.4: Specifications of computational resources used for demonstration model

	PC	HPC	GCP
Max. # cores	8	28	64 (tested up to 32)
CPU speed	3.60 GHz	2.4 GHz	2.0 GHz
Processor type	Intel i7	Intel Xeon	Intel Xeon
RAM	16 GB	128 GB	7-120 GB (depending on # of CPUs)

4.3 Results and Discussion

4.3.1 Results from Case A and Case B

Figure 4.7 shows the results from the three control scenarios for Case A. In the rules-based and MPC scenarios, the control policies kept the valves closed more, thus retaining more water in the storage units and preventing flooding at Node J3 (see Figures 4.7 A and B). The water level at St1 reaches much higher values in Scenarios 2 and 3 (max of 1.45 m (4.76 ft) and 1.38 m (4.54 ft)) compared to Scenario 1 (max of 1.26 m (4.15 ft)). In this case, both rules-based control and MPC practically eliminated flooding while in the passive scenario, flooding occurred (Figure 4.7D).

In addition to practically eliminating flooding in Case A, the swmm_mpc control policy was able to maintain the water levels in St1 and St2 near the target of 1.22 m (4.00 ft). Because the cost of flooding was weighted more heavily than the cost of deviating from target water levels at St1 and St2, the storage units allowed the water levels to exceed the target depth. Following the peak of the storm, and therefore the largest risk of flooding, the algorithm could focus on the less-weightier objective of minimizing deviations from the target water levels. Therefore, following the peak of the storm, water was released from the storage units until the level reached near the target water level. The final water levels of St1 and St2 were 1.20 m (3.94 ft) and 1.17 m (3.83 ft), respectively.

Figure 4.8 shows the results for Case B, in which less storage volume is available than in Case A. In Case B, the rules that completely eliminated flooding in Case A were not effective and in fact resulted in more flooding than the passive scenario. In both Case A and 2, the flooding in the passive scenario occurred at Node J3, downstream from the storage units. By contrast, in Case B the flooding in Scenario 2 occurred at the storage units. In this instance, the rules held back too much water so when the peak of the storm arrived, the storage units overtopped. As with Case A, the policy found by swmm_mpc, was effective at practically eliminating flooding. This illustrates that there are conditions in which one set of simple rules is effective at achieving an objective and other conditions in which it is not. In Case B, with less storage available, the more sophisticated, more computationally expensive control policy found through swmm_mpc was needed to achieve the objective of minimizing flooding.

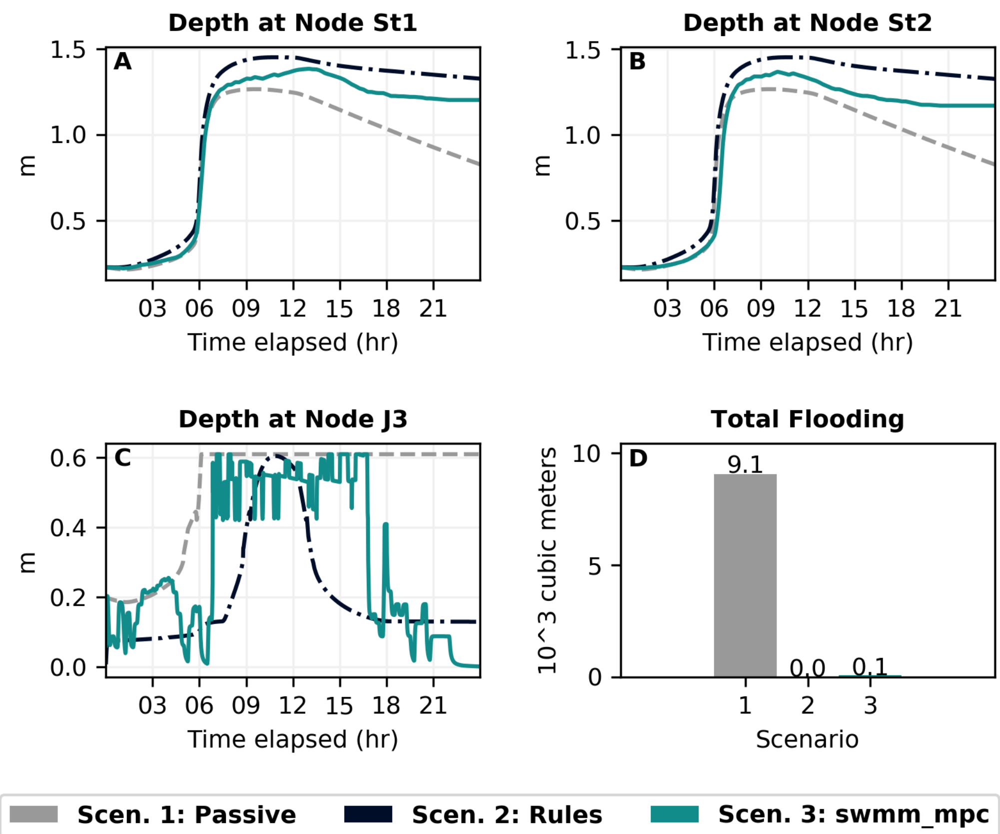

Figure 4.7: Depth and flooding at system nodes for Case A

In both Case A and Case B, a major difference between the rules-based and the swmm_mpc scenarios is the smoothness that is seen in the control policy. The swmm_mpc control policies changes the valve position and therefore the upstream and downstream water depths much more abruptly than the rules-based policies. This abruptness could be smoother if the time for the orifices to implement a change in setting was increased. In my demonstation model these times were zero meaning that the orifice settings were changed instantly.

4.3.2 Computational cost of swmm_mpc

Figure 4.9 shows the wall-clock times for executing swmm_mpc for Case A on a PC, an HPC, and GCP machines with a varying number of processing cores. The simulation had 96 control time steps (15 minute resolution for 24 hours). If used for online MPC in a real case, the wall-clock time required for one time step would need to be less than the time step itself, otherwise, the setting for the next time step would not be determined before it would need to be implemented.

The fastest wall-clock time using the PC was 89.4 minutes using eight computational cores (the maximum available). Therefore the time required to find the best control policy at each control time step was 0.93

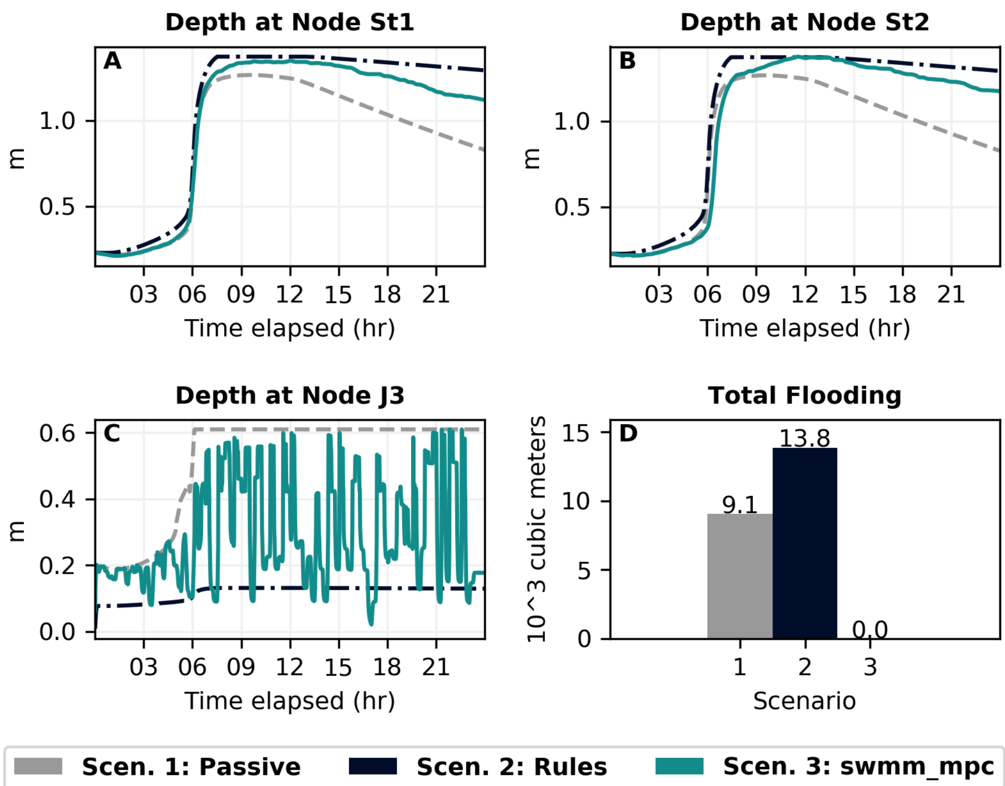

Figure 4.8: Depth and flooding at system nodes for Case B

minutes. In this case, the PC's computing power was sufficient (0.93 minutes per time step compared to 15 minute time step). For the HPC, the best case scenario was a wall-clock time of 18.2 minutes (0.19 minutes per time step) using the maximum of 28 computational cores. Although the minimum wall-clock time was achieved using all 28 cores on the HPC, the improvement in wall-clock time when increasing the number of cores past 16 was minimal. This is a relevant consideration when using a shared HPC resource where requesting more computational cores likely corresponds to a longer wait in the job queue.

The wall-clock times using GCP were much lower than the PC or HPC. In the best case, 32 vCPUs and 120 GB of RAM were used for a wall-clock time of 7.47 minutes (0.08 minutes per time step). This is a 2.4x speed-up compared to the fastest run using the HPC and almost a 12x speed-up compared to the fastest PC run. The financial cost of renting this machine was $1.71 per hour. The GCP hardware is newer which may explain why the wall-clock time is lower even when the same number of computational cores was used.

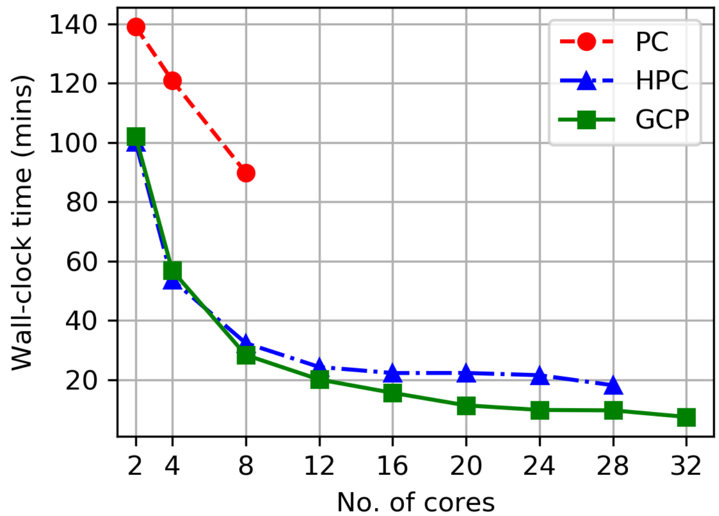

Figure 4.9: Wall-clock run times for Case A with varying number of computational cores using a PC and a high performance computer.

4.3.3 Pratical considerations

Computational costs

The execution times for my example use case were viable, however, a more complex model, a smaller control time step, or different EA parameters (more generations or individuals per generation) would increase the execution time. For example, the demonstration SWMM5 model required only one second or less to execute. The swmm_mpc workflow executed that model thousands of times. In my cases, the model was run more than 30,000 times (24-hour simulation x 4 time steps per hour x 5 generations per time step x approx. 70 individuals per generation) therefore requiring approximately 30,000 seconds of computation time (if each model takes approx. 1 second to run). If a more sophisticated SWMM5 model instance were used, the execution time would be much higher. For example, in related research I am using a more complex model of the stormwater infrastructure for a neighborhood in Norfolk, Virginia that requires close to 60 seconds to execute for a 24-hour simulation time period. The wall-clock time for swmm_mpc for this more complex model would increase by around a factor of 60 compared to the simple cases demonstrated here. Assuming a linear increase, the wall-clock time would be 55.8 minutes/time step using the PC, thus rendering it unfeasible for running on a PC. Again assuming linear scaling, using 32 cores on GCP, the same simulation would execute at a rate of 4.8 minutes/time step. Using this setup, the wall-clock time for a 24-hour simulation with a time step of 15 minutes would be approximately 7.68 hours.

Another factor to consider for practical use of swmm_mpc is the control horizon and the number of control structures whose policies will be found using swmm_mpc. These two parameters determine the size of the overall control policy and therefore the solution space that the EA will be searching. In my example use case, the control policy was a vector of eight integers between 0 and 10. Therefore, there were 11^8 possible solutions. This solution space, already large, would double if the control time step were 7.5 minutes instead of 15, or if the control horizon were two hours instead of one. A larger solution space would result in a larger computation time to reach an effective solution.

Given the computational cost of the current swmm_mpc approach, the required complexity (and thus the wall-clock time) of the process model is an important consideration. A scenario with a simple process model (approx. 1 second wall-clock time) can be feasibly executed with just a PC, as shown in the system demonstration. However, a simple model may not represent complex urban drainage systems well enough to produce an effective control policy. Defining what level of detail is sufficient in the process model may be difficult, however, as it may depend on the objective of the modeling, a certain storm event, or the system itself. There is a tradeoff among 1) model complexity, 2) model runtime, and 3) the model's ability to effectively simulate the relevant parts of the system. This tradeoff is very relevant to the use of process models in a receding control horizon approach such as MPC and needs further research.

If a more complex model is needed, municipalities or others needing cloud-based resources to run swmm_mpc must consider the financial cost of renting a machine. Using GCP, the cost of finding the control policy for the 24-hour time span in Case A in my system demonstration was very low, $0.21. This was, however, for one simple case. The cost would be higher with more complex scenarios such as a more complex model, a shorter time step, or more controls. If I assume the use of a more complex model, which takes 60 seconds to run, would increase running time by a factor of 60, that would also increase the cost by a factor of 60 to $12.60. However, the system would only run if there were a storm in the forecast.

Given the computational cost of running the evolutionary algorithm, other, more efficient alternatives should be explored in future research. One possible alternative that is reinforcement learning (Kaelbling et al., 1996). This approach may be able to converge to a solution more quickly than an evolutionary algorithm and thus reduce runtimes. Another future improvement could be adding a penalty to changing actuator states and/or using another dynamic optimizer to have a less erratic behavior in the actuators.

Data and modeling uncertainty

The current design of swmm_mpc does not take into account the uncertainties in the system states,the forecast data and the process model. Because the SWMM5 engine is used to simulate both the urban drainage system and the process model, the process model assumes 1) perfect knowledge of the urban drainage system states, 2) perfect knowledge of future disturbances, and 3) perfect modeling of the urban drainage system. In a real implementation, there would be significant uncertainties in each of these aspects. In a real implementation, knowledge of the system states is available only from a limited number of sensors in the system. This data, limited in spatial and temporal resolution, would need to be interpolated, and likely extrapolated, to set all the states in the system. More work will need to be done to investigate ways of incorporating sensor values to set the process model's initial conditions. Additionally, in the current case, the future disturbances (i.e., primarily rainfall) are known perfectly, when in reality, there is a large amount of uncertainty involved with forecasting such disturbances (see for example, Hong and Pai (2007), Valverde Ramírez et al. (2005), and Bellon and Austin (1984) regarding uncertainty in forecasting rainfall).

In addition to data uncertainties seen in reality, swmm_mpc does not currently consider gaps between the *simulated* behavior through the SWMM5 process model and the *actual* behavior of the urban drainage system, but assumes that simulation and reality are the same. In actuality, the gap between simulation and reality in urban drainage systems can be significant (see for example Mark et al. (2004b)). On a related note, the ability for swmm_mpc to find a control policy that is effective for the urban drainage system is directly related to how well the process model represents the system. Given the simulation and data gaps seen in reality, the simulated results through policies found by swmm_mpc should be considered as the best case scenario and if the same policies were used in practice, any effects should be expected to be seen to a lesser extent. Further research is needed to determine the degree to which the results from the policies implemented in reality will differ compared to the simulation results.

4.4 Conclusions

A free and open-source software package, swmm_mpc, was developed which computes a control policy for controls within an urban drainage system model. The widely-used United States Environmental Protection Agency Stormwater Management Model Version 5 (SWMM5) is used to simulate the urban drainage system and the process model. A third-party Python library, pyswmm, is a critical component of the swmm_mpc workflow allowing a SWMM5 model to be run step-by-step in a Python environment. An evolutionary algorithm was used to find an effective control policy at each time step. When tested using a simple SWMM5 model, the swmm_mpc software was able to produce control policies that met objectives including minimizing flooding and minimizing deviation from target water levels at certain nodes in the system.

swmm_mpc leverages parallel computing to run the computationally expensive evolutionary algorithm more quickly. The wall-clock time for a simple SWMM5 model for a 24-hour simulation was reduced from 139 minutes to 89.4 minutes when the computational cores on a desktop PC were increased from two to eight. The wall-clock time was reduced even further to 18.2 minutes on a 28-core high-performance computer and to 7.47 minutes on a 32-core machine rented through the Google Cloud Platform. Parallel computing makes swmm_mpc feasible for use in real-time control with complex process models.

As the average storm intensity is projected to increase, and sea levels are expected to continue to rise, cities globally and especially on the coasts, can expect more flood conditions. Active control of urban drainage systems may be one of an array of approaches that can be used to confronting these challenges. The swmm_mpc software I have developed can be used, built-from, and improved upon as a tool to assist decision-makers and researchers in finding effective control policies for urban drainage systems.

4.5 Software Availability

The swmm_mpc software is open-source and available for use and improvement on GitHub at https://github.com/UVAdMIST/swmm_mpc. A Docker image of swmm_mpc is also available at https://hub.docker.com/r/jsadler2/swmm_mpc/. The demonstration model is available on HydroShare (Sadler et al., 2018b).

4.6 Acknowledgments

I also acknowledge the support of awards 1735587 (Critical and Resilient Infrastructure Systems - CRISP) and 1737432 (Smart and Connected Communities - S&CC) from the United States National Science Foundation. I thank Bryant McDonnell, one of the main developers of pyswmm, for his assistance in adding the hotstart file functionality to OWA-SWMM5 and pyswmm. I also thank the consultants of the University of Virginia Graduate Writing Lab including Kelly Cunningham, its director, for their helpful feedback in preparing the manuscript.

Chapter 5

Utility of Model Predictive Control of Stormwater Systems with Sea Level Rise

5.1 Introduction

5.1.1 Increasing flooding in coastal cities with SLR

Coastal urban flooding is likely to increase globally in the coming decades. In addition to more intense rainfall (Berggren et al., 2012; Neumann et al., 2015), which can affect any city, expected sea level rise (SLR) (Church et al., 2013) makes coastal cities particularly vulnerable to increased flooding. In coastal cities, the water level of the receiving water body has a large impact on drainage and flooding. When receiving water levels are above the system outlets, system pipes cannot drain and when this happens during a storm event, backups can occur. In such cases, not only is there a decrease in the amount of volume available in the system, but the hydraulic head is also decreased, slowing the rate of drainage. If the receiving water levels are high enough, water can back up into stormwater pipes and flows back through stormwater inlets in the city, causing flooding even with no rainfall at all. SLR will make these problems worse as the average water level at system outfalls increases.

In addition to the magnitude of tide cycle, the timing of the tide cycle can significantly influence flooding in coastal cities. For coastal cities, the stormwater outfalls could be inundated at high tide and completely free at low tide. Thus, the ability for the stormwater system to drain runoff from the city may be affected by the timing of the storm's runoff relative to the timing of the tide cycle.

5.1.2 Real-time controls of stormwater systems

One strategy for mitigating flooding in coastal cities is the use of real time control (RTC). RTC of stormwater systems consists of three major components: (1) real-time sensors, (2) analysis, and (3) actuators. Sensors (e.g., rain gauges, level sensors, flow meters) provide real-time information on system states. Given the real-time system states, a control decision is made based on either offline or online system analysis. The control decision determined upon in the analysis is then implemented by actuators (e.g., pumps, gates, valves) which affect the behavior of the system until the next control decision (Schütze et al., 2004).

Although conventional stormwater systems are driven by gravity and behave statically, RTC system actuators counteract gravity to control the storage and flow of water dynamically. Using RTC, a stormwater system can be managed in a more optimal way (Kerkez et al., 2016). For example, there may be under-utilized capacity in the system in the form of unused storage in a pond or inline that could be utilized with RTC devices. If the existing capacity is used more effectively, there is less need for expansions of physical infrastructure.

RTC in urban drainage systems has been primarily explored, implemented, and evaluated for combined sewer systems (CSS) with the primary objective of minimizing combined sewer overflows (CSOs) (Meneses et al., 2018; Garofalo et al., 2017; Pleau et al., 2005; Schütze et al., 2004). Although the research on RTC for CSS has been well-developed in the literature, there has been a smaller but growing amount of research on the use and utility of RTC in separate sewer systems (Wong and Kerkez, 2018) which are used in most modern cities. Additionally, to my knowledge, there has been no research about how RTC systems could benefit coastal cities given their current and projected flooding challenges due to SLR. It is therefore the intent of this paper to characterize the utility of RTC, specifically an implementation of model predictive control (MPC) (Puig et al., 2009; Cembrano et al., 2004; Schütze et al., 2004; Gelormino and Ricker, 1994) in a separated sewer system subjected to various sea level rise scenarios.

While the use of RTC systems could benefit coastal cities globally, the geographic focus of my study is the coastal city of Norfolk, Virginia USA. Norfolk is part of a region which is experiencing faster than average SLR due to land subsidence (Mitchell et al., 2013). Therefore, the conditions being seen now in Norfolk and its neighboring cities could be seen in coastal cities globally in the decades to come.

We will assess the current and future utility of MPC systems in Norfolk using a simulation model, the Environmental Protection Agency (EPA) Stormwater Management Model version 5 (SWMM5). In the SWMM5 model, three scenarios will be simulated with increasing sea levels: 1) the passive scenario (the current state of the system), 2) the passive scenario with the addition of a tide gate (a mechanical device used to prevent backflow into the system), and 3) the tide gate scenario with the addition of actuators controlled by MPC. The simulated flood volumes from the three system scenarios will be compared and flood reduction from MPC calculated. The increasing receiving water levels will then be related to SLR scenarios. In addition to increasing the water level at the stormwater outfalls, I will also vary the timing of the tide cycle and assess the effect of this factor on flooding.

In the remainder of this paper, I present details about my methods including the study area of Norfolk, VA, the MPC scenarios explored, and the rainfall and tidal conditions used. I then present the results which demonstrate the utility of MPC (in terms of flooding) with increasing sea levels. Finally I discuss the results and relate the increasing receiving water levels to potential SLR scenarios.

5.2 Methods

5.2.1 Study area

Norfolk, VA USA served as the study area for this research (Figure 5.1). Norfolk is a peninsula with several inland streams and much of the city is within close proximity to a tidally-influenced water body. The city is low-lying with low topographic relief. The average elevation of the city is less than 10 feet. Because of these geographic conditions, Norfolk experiences frequent flooding (Sadler et al., 2018a).

Given the size of the city and the complexity of the stormwater system with all its features (subcatchments, pipes, channels, junctions, ponds, etc.), my study focused on a subsection of the city: the neighborhood called the Hague. The Hague, approximately eight square km, is a key part of Norfolk. It is home to many of Norfolk's most historic buildings and cultural attractions. The Hague is also key to the city's connectivity as it is adjacent to the city government buildings and the regional hospital. In addition to its importance to Norfolk, the Hague area is one of the most flood-prone of the city.

5.2.2 Storm event and base tidal conditions

There are two main factors that affect flooding in low-lying coastal cities: rainfall and tidal conditions. Because this paper is primarily focused on the utility of MPC over time given SLR, the rainfall event in all the simulations was kept constant and the tidal conditions were varied. The rainfall event I used was a 2-year 12-hour storm event (3.08 in) (Bonnin et al., 2018) with a SCS Type II temporal distribution (Mockus, 2012) (Figure 5.2).

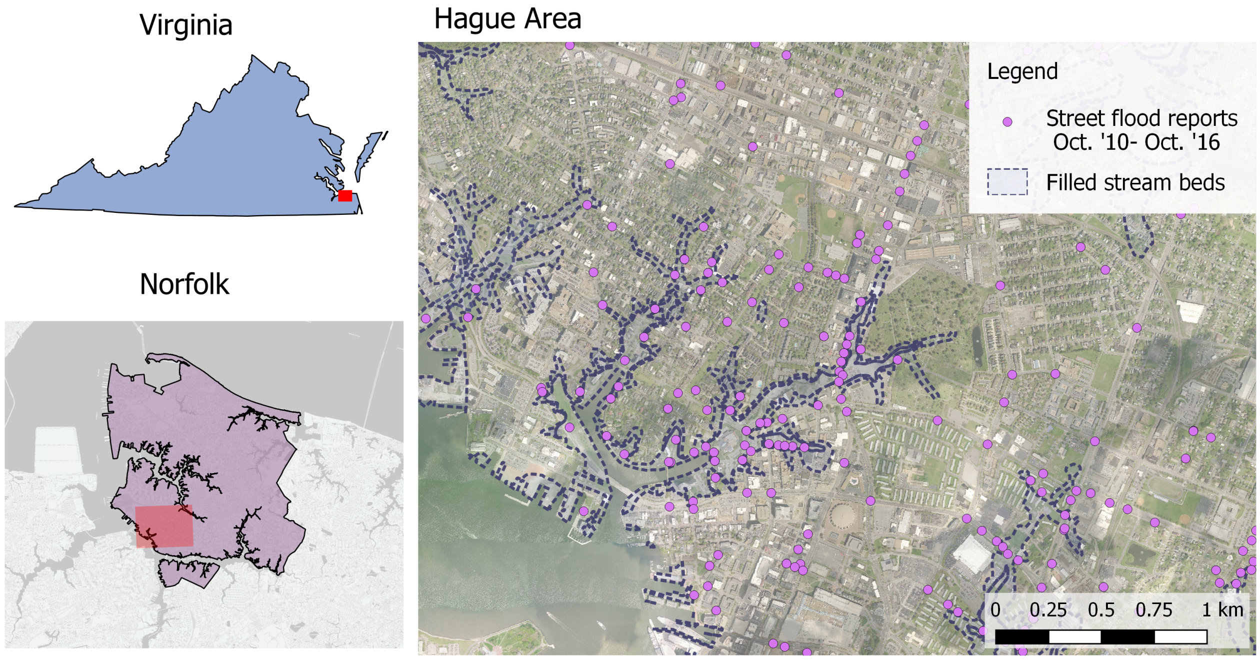

Figure 5.1: Study area - Hague neighborhood in Norfolk, VA USA

Observations from the Sewell's Point tide gauging station, operated and maintained by the US National Oceanic and Atmospheric Adminstration (NOAA) (NOAA, 2017e), served as the baseline tide conditions in the simulations. Hourly water level observations from a typical 24-hour tide cycle were taken as a base tidal boundary condition in the simulations. The tide cycle (see Figure 5.2) was the hourly recorded water level recorded at the Sewell's Point station on August 8, 2018. This day was selected because the 24-hour tidal range this day (2.78 ft) was very similar to the station's all-time average 24-hour tidal range (2.76 ft) and because the day had a very typical tide cycle without unusual effects from wind or other factors.

5.2.3 Increasing tide levels

To estimate the utility of MPC with SLR, the baseline hourly tidal conditions were increased in magnitude. To cover a range of possible SLR scenarios, the hourly tide levels were increased by 0.5 foot increments up to 3.5 ft, for a total of seven elevated water levels. The increases in tide level were correlated to region-specific SLR estimates which took into account local land subsidence (Mitchell et al., 2013) .

5.2.4 Computational model

The computational model used to simulate the storwater system in the study area was the Environmental Protection Agency Stormwater Management Model version 5 (EPA-SWMM5). SWMM5 was used based on

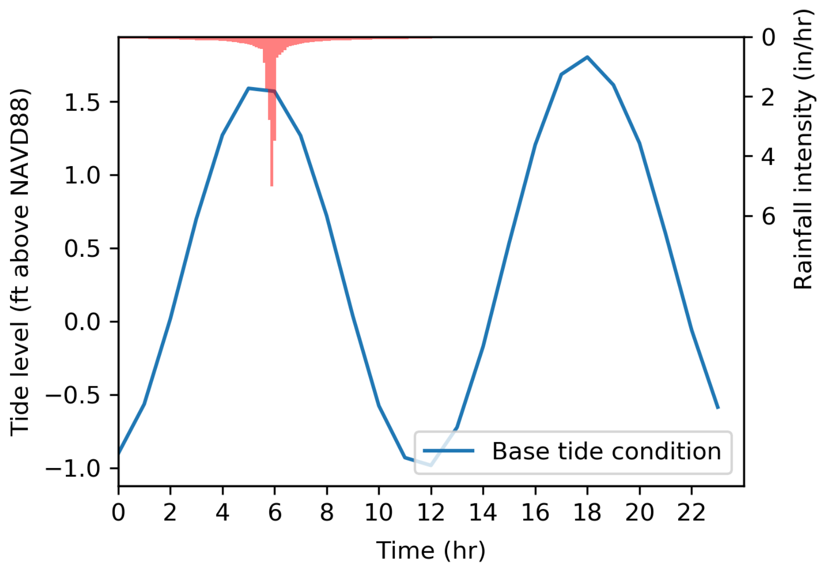

Figure 5.2: Design storm and base tide condition

its general acceptance and use as a 1D stormwater system model (Burger et al., 2014). SWMM5 solves the St. Venant equations numerically to simulate dynamic routing in a stormwater system. The use of dynamic wave routing is particularly important in simulating coastal systems because the systems can experience backflow from tidally influenced receiving waters. In a SWMM5 simulation, each node has a maximum depth (usually the distance from the invert elevation to the ground surface elevation). When the simulated water exceeds the maximum depth at a node, the excess water volume is recorded as flood volume by SWMM5. At the end of a simulation, both the total amount of flood volume at each node and all the nodes together is reported.

A SWMM5 model of the Hague neighborhood was obtained from the city of Norfolk (Figure 5.3). The city of Norfolk verified that the model behavior the behavior of the physical system. The model consists of 208 nodes with a wall-clock time of approximately 0.5 minutes for running a 24-hour simulation of a 12-hour storm event on a 8 core desktop computer. The system drains to a tidally influenced water body and a water level boundary condition was implemented at the outfall to simulate the influence of the tidal water body at the receiving waters.

5.2.5 Addition of tide gate

Before simulating MPC in the passive system, a tide gate was added to the outfall of the SWMM5 model. A tide gate is a passive device used to prevent backflow from a receiving water body. The use of a tide gate was explored because with SLR, the likelihood of backflow increases and the installation of a tide gate could become an important approach for minimizing tidal flooding in coastal cities. In SWMM5 simulations, tide gates are assumed to be 100% effective at preventing backflow through the system outfall. Additionally, in a SWMM5 simulation, water cannot flow out of an outfall with a tide gate when the outfall is inundated.

5.2.6 Addition of MPC

Three actuators for mitigating flooding in coastal cities were simulated in the MPC scenario. First, a valve at the outlet of a simulated pond (Point A in Figure 5.3) was used to control the storage volume in the pond. Second, a pump (Point B in Figure 5.3) was simulated to increase the hydraulic head in the stormwater pipes. This is particularly important in a low-relief, coastal city because the contribution of elevation head to the total hydraulic head is very small. Third and finally, an inflatable dam was simulated to utilize inline storage in the stormwater pipes (Point C in Figure 5.3). Although these controls were simulated in the Hague area, such controls (or different combinations of controls) would likely be useful in many coastal cities globally since many coastal cities share similar geographic characteristics (low relief, low elevation).

The control approach in the MPC scenario was simulated using the swmm_mpc Python library (see Chapter 4). The main parameters in running swmm_mpc include cost function parameters, genetic algorithm parameters, and control parameters. The cost function in these simulations was the total amount of flooding that occurred in the simulation as reported in the SWMM5 output. For the genetic algorithm, six generations were used (including the initial generation) with twenty individuals in the initial generation. The control horizon was 30 minutes and the control time step was 15 minutes, meaning that the algorithm was searching for two control settings for each actuator at every time step in the model (see Table 5.1).

In MPC, the control strategy for each actuator is determined and can be changed at every control time step (15 minutes in this case). This makes the approach flexible, that flexibility, however, comes at a higher computationally cost since the control strategy is determined by running many iterations of the simulation model at each time step. Given the computational cost of running this library, the MPC scenarios were executed on a high-performance computer with 20 computational cores and 128 GB of RAM.

Table 5.1: Control parameters for MPC scenarios

Parameter	Value
Control Horizon (hr)	0.5
Control Time Step (hr)	0.25
Num controls	3

Assessment of MPC with SLR

To assess the utility of adding MPC in the Hague neighborhood stormwater system, the storm event and tidal conditions, described above were input into swmm_mpc. For each of the tidal and gate scenarios, the amount

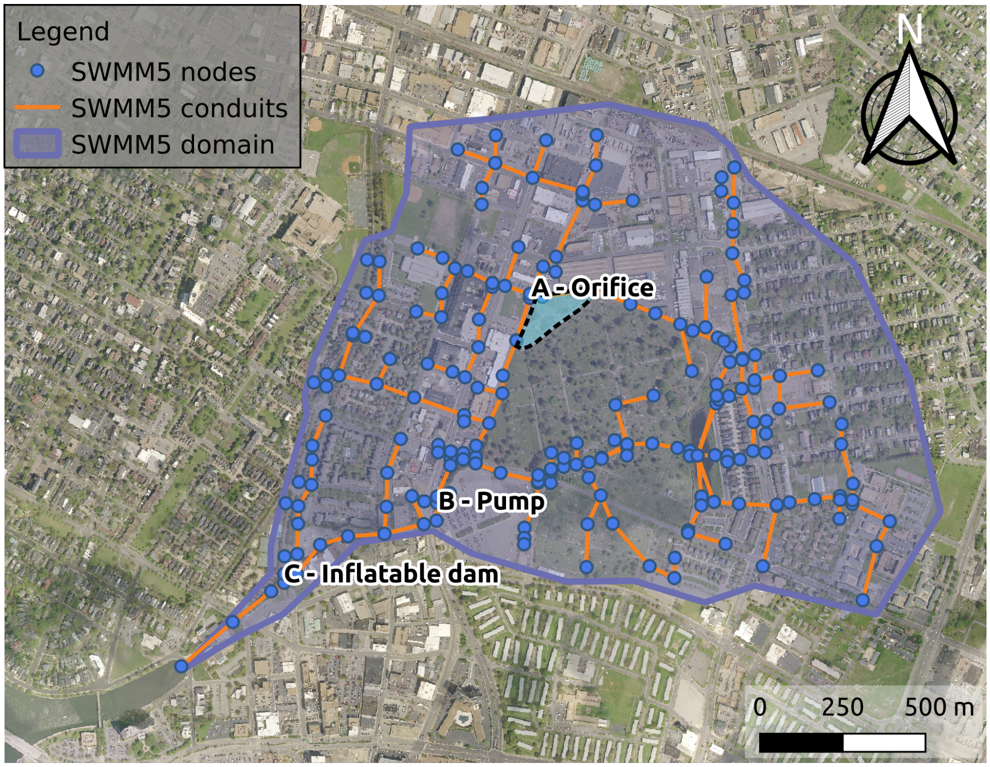

Figure 5.3: SWMM5 model of the study area and simulated actuators

of flooding that occurred in the MPC scenario was compared to the amount of flooding that occurred in the scenario with the tide gate. The percent reduction of total flooding and the maximum flooding at a single node between the MPC and the tide gate scenarios was also calculated. In preliminary simulations, it was found that the runoff from the rainfall event alone caused flooding at some nodes that were not affected at all by tidal conditions or MPC. The majority of this flooding occurred upstream of the controls at three nodes. Therefore, this volume of flooding, which was in total 0.582x10^6 gallons, was not taken into account when calculating the percent decrease in total flooding. Instead, an effective percent reduction of total flooding was calculated:

$$\text{effective percent reduction} = \frac{F_{G\_eff} - F_{R\_eff}}{F_{P\_eff}} \tag{5.1}$$

and

$$F_{G\_eff} = F_G - F_{uneff} \tag{5.2}$$

$$F_{R\_eff} = F_R - F_{uneff} \tag{5.3}$$

where F_G and F_R are the total simulated flood volumes in the tide gate and MPC scenario, respectively, and F_{uneff} is the flood volume unaffected by MPC and tidal conditions (0.582x10^6 gallons in this case).

5.2.7 Timing of tide cycle

In addition to the magnitude of the tide levels, the timing of the tide cycle relative to the timing of a rainfall event can impact drainage and flooding in a coastal stormwater system. Therefore, in addition to analyzing changes in flooding with increasing the magnitude of tide levels (the main focus of this study), the timing of the tide cycle relative to the 2-year 12-hour design storm was also varied. To quantify the effect of this factor, the timing of the design rainfall event was held constant while base tide cycle was changed was shifted in 1 hour increments from 1-10 hours. The amount of flooding that occurred with each shift in timing of the tide cycle was recorded. This was done under increases in sea level as above (up to 3.5 ft increase in 0.5 ft increments). The tide cycle timing which caused the most flooding for the most increases in sea level was used for the remainder of the analysis. Because the timing of the tide cycle would affect both the MPC and passive systems similarly, the different timings were only simulated on the passive case.

5.3 Results

5.3.1 Flooding with SLR in passive scenario

Figure 5.4 shows the amount of flooding in the passive system given increases in sea level. The amount of flooding remains fairly constant between an increase in sea level of 0 ft 1.5 ft, however, there is a dramatic increase in flooding between increases of 1.5 ft and 2.0 ft. This occurs because with an increase of 2.0 ft, the high tide (3.8 ft) exceeds the elevation of lowest maximum depth elevation in the system, 3.41 ft at node E143274 (see Figure 5.5). When the tide exceeds 3.41 ft, the water from the receiving water exits node E143274 and causing flooding. This occurs when the base tide cycle is increased in magnitude by 1.6 ft. A sea level increase of 1.6 ft therefore becomes a tipping point and any increase in sea level beyond this causes much more flooding. The flood volume increases substantially with each increase above 1.5 ft.

5.3.2 Flood reduction with addition of tide gate

The impact of adding a tide gate to the outfall of the system dramatically reduced flooding with increasing sea levels compared to the passive system (Figure 5.6). This reduction was most evident with sea level increases of 2.0 ft and above when the vast majority of the flooding was tidally driven. With smaller increases in sea level, the effect was less significant. This is most striking between increases of 1.5 ft and 2.0 ft. The

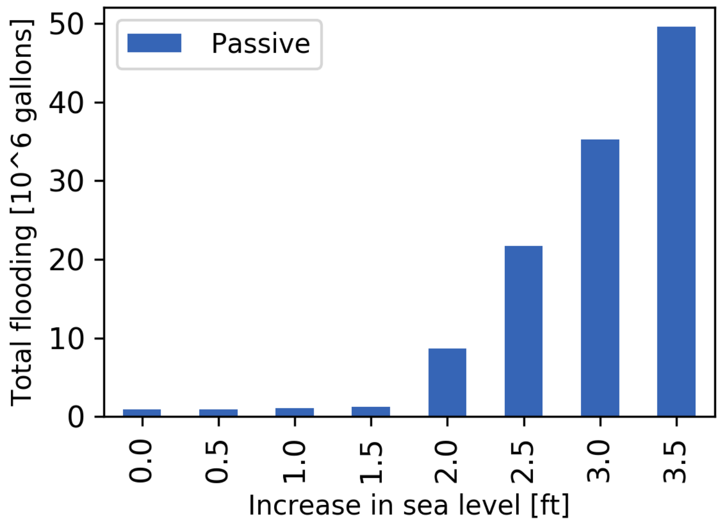

Figure 5.4: Flooding increases with sea level rise

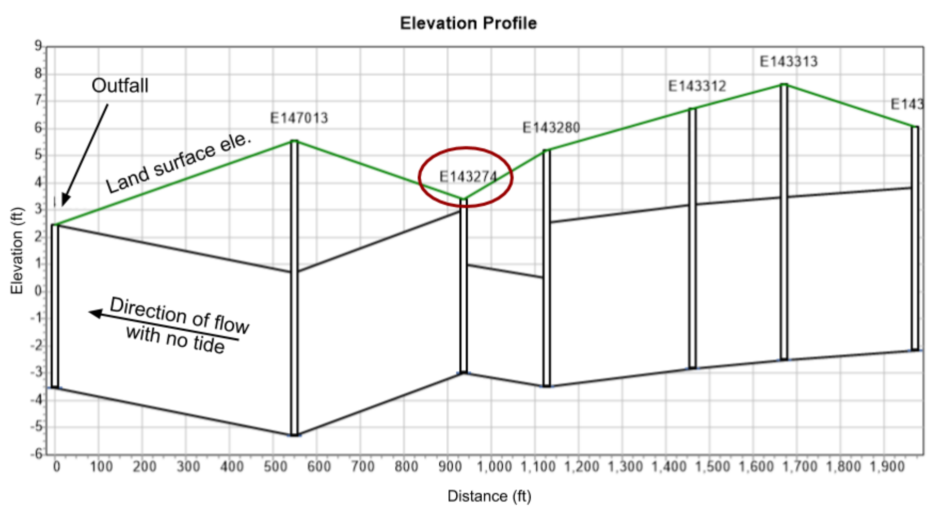

Figure 5.5: Elevation profile of the closest nodes to the drainage outfall

addition of the tide gate with a 1.5 ft increase in sea level reduced total flooding from 1.3x10^6 gallons to 1.1x10^6 gallons, a reduction of 15%. In contrast, with a sea level increase of 2.0 ft, the addition of a tide gate reduced the total flood volume from 8.7x10^6 gallons to 1.7x10^6 gallons , an 80% reduction. When sea level was increased to 3.5 ft, flooding was reduced from 49.6x10^6 gallons to 3.3x10^6 gallons, a 93% reduction.

5.3.3 Flood reduction with tide gate and MPC

Figure 5.7 shows the amount of flooding seen in the tide gate and MPC scenarios given increases in sea level. Like in the tide gate scenario, the total amount of flooding increases with increasing sea level. However the

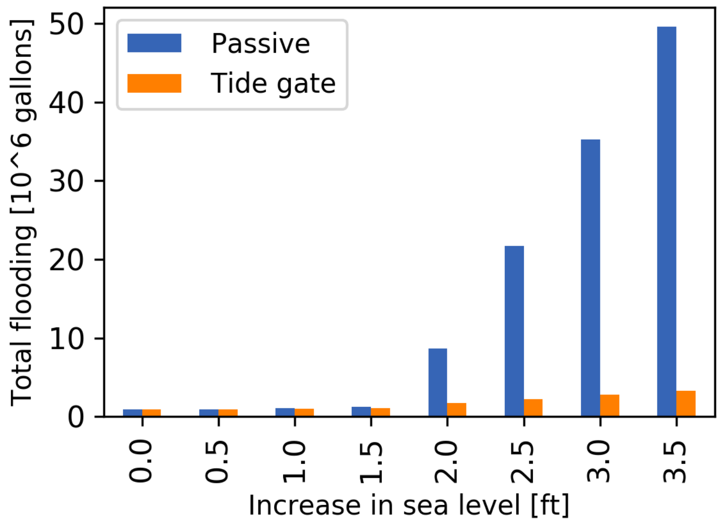

Figure 5.6: Flooding increases with Sea Level Rise

rate of increase is slowed as the actuators in the MPC are able to reduce flooding at a rate higher than the increase in flooding in the situation without MPC (see Figure 5.8). For example, from a 3 ft to a 3.5 ft increase in sea level, the flooding increases by 18% without MPC; with MPC the increase is only 9%.

MPC reduced flooding by at least 40% for each of the increased sea levels below 2.0 ft (Figure 5.9). Above 2.0 ft increases in sea level, the percent reduction generally increased with increases in sea level up to a 35% reduction in flooding with a sea level increase of 3.5 ft. In addition to reducing the total flood volume, MPC reduced the maximum node flood volume with increasing sea level after an increase of 2.0 ft (Figure 5.9). The largest reduction of maximum node flood volumes was 70% at Node E143274 with a 3.5 ft sea level increase (from 2.2x10^6 gallons without MPC to 0.65x10^6 gallons with MPC).

5.3.4 Timing of tide cycle

In a similar trend to the increase in flooding with SLR, the difference between the maximum and minimum flood volumes from different timings also increases substantially starting with an sea level increase of 2.0 ft (Figure 5.11). The difference between the maximum and minimum flood volumes, however, does not follow the same trend when divided by the average flood volume at each sea level increase (Figure 5.12). The peak normalized difference, in fact, occurs at a sea level rise increase of 1.5 ft. These results suggest that following a sea level increase of 2.0 ft, the tidal flood volume is so large that the flood volume from the rainfall event, and thus the timing of the tide cycle relative to the rainfall, makes increasingly smaller impact.

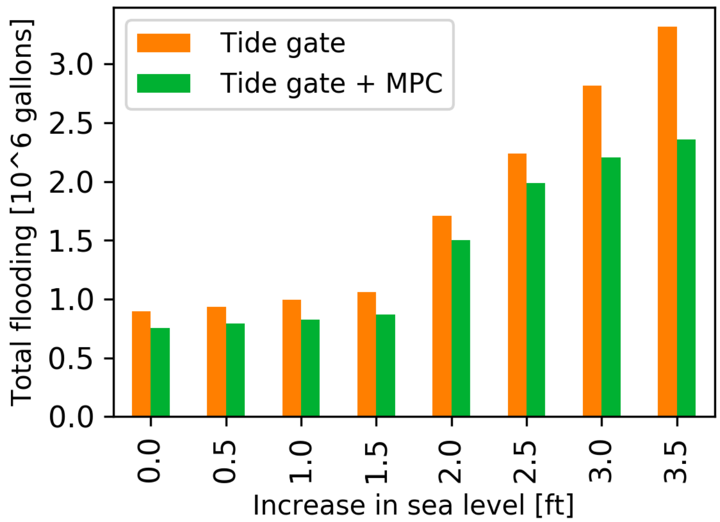

Figure 5.7: Amount of flooding of tide gate scenario and MPC scenario given increasing sea level

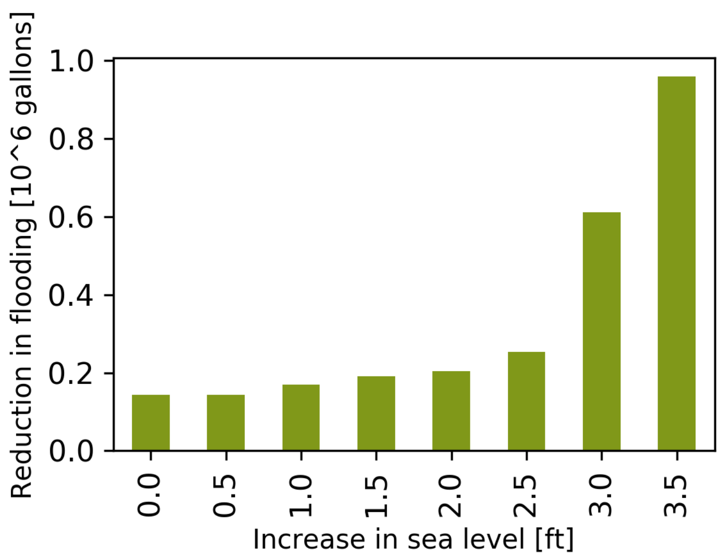

Figure 5.8: Flood reduction in MPC compared to tide gate scenario given increasing sea level

5.4 Discussion

Although the use of a tide gate makes a larger difference in overall flooding compared to the passive scenario, the use of MPC further improves the system performance. Not only does MPC reduce the overall volume of flooding, but it also reduces the maximum node flood volume, thus more evenly distributing the flood volume across the nodes in the system. Less severe flooding at any given node could mean that less disruption and or damage is occurring.

In my analysis, the utility of MPC, in terms of percent flood volume reduction and absolute flood volume

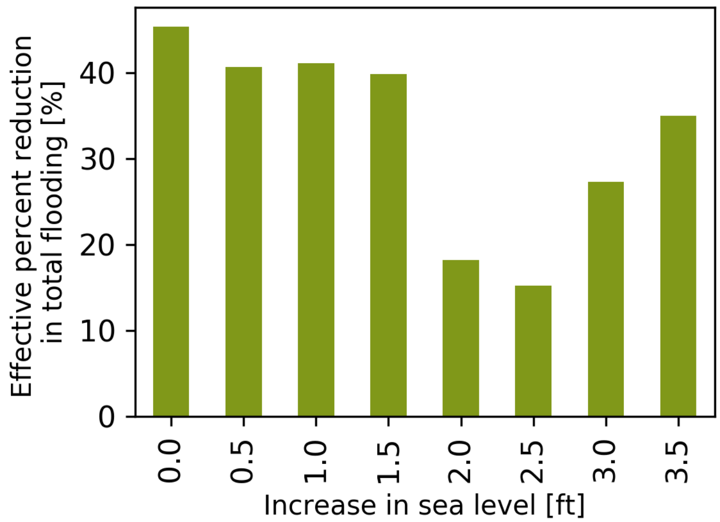

Figure 5.9: Effective percent total flood reduction in MPC compared to tide gate scenario given increasing sea level

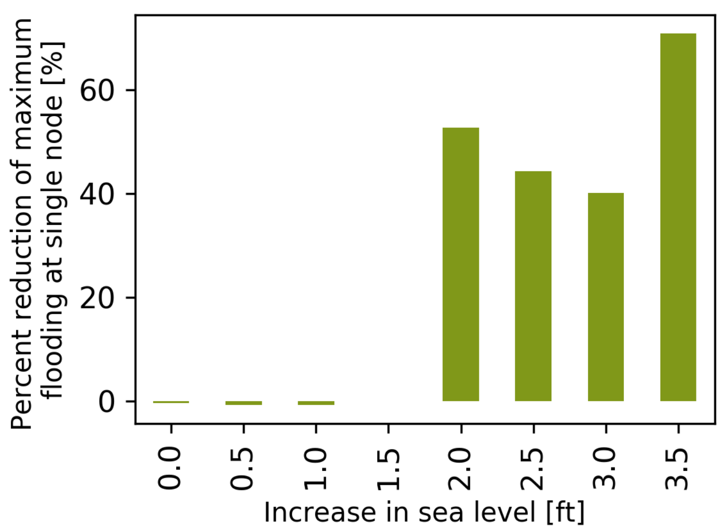

Figure 5.10: Percent reduction of maximum flooding at an individual node with MPC compared to tide gate scenario given increasing sea level

reduction, generally increased with increasing sea levels. The largest increases in utility occurred with higher increases (3.0 ft and 3.5 ft). This suggests that as coastal cities consider investment and mitigation alternatives, the investment in RTC (and particularly MPC) could provide a greater return on investment as time passes.

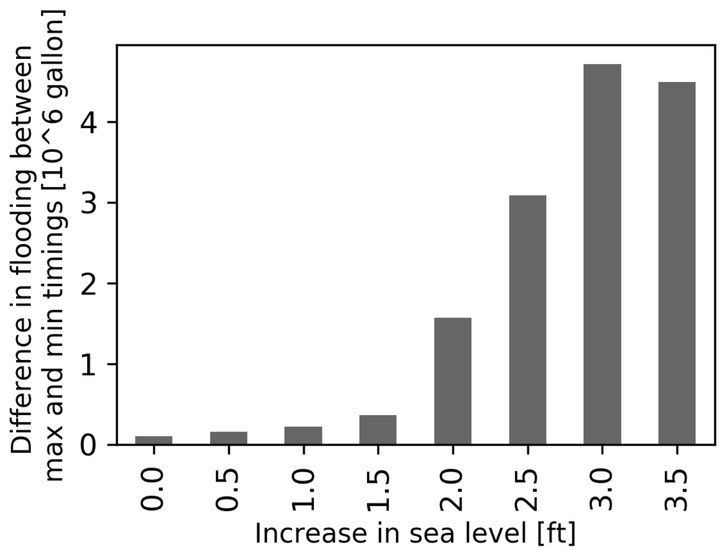

Figure 5.11: Difference in flooding between timing that caused the most flooding and timing that caused the least for each increase in sea level

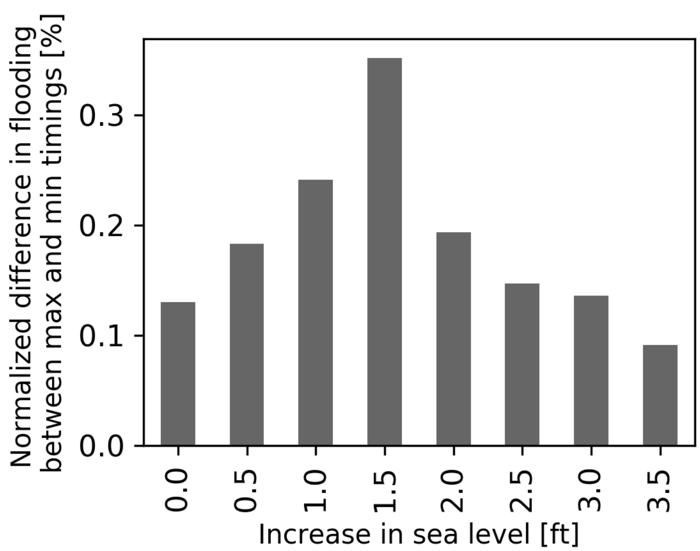

Figure 5.12: Difference in flooding between timing that caused the most flooding and timing that caused the least for each increase in sea level

5.4.1 Correlation with SLR and Extreme Water Levels

Figure 5.13 shows SLR estimates specific to south-east Virginia (Mitchell et al., 2013). The "high", "medium" and "low" scenarios are different emission scenarios. The "historic" is an extrapolation of the existing tidal record. As sea level increases with time, a normal tide cycle will become increasingly problematic for coastal cities. Based on my analysis above, there will be a dramatic increase in tidal flooding once the high tide reaches 3.41 ft. This would be an increase of the base tide cycle in Figure 5.2 of 1.6 ft (shown as the dashed

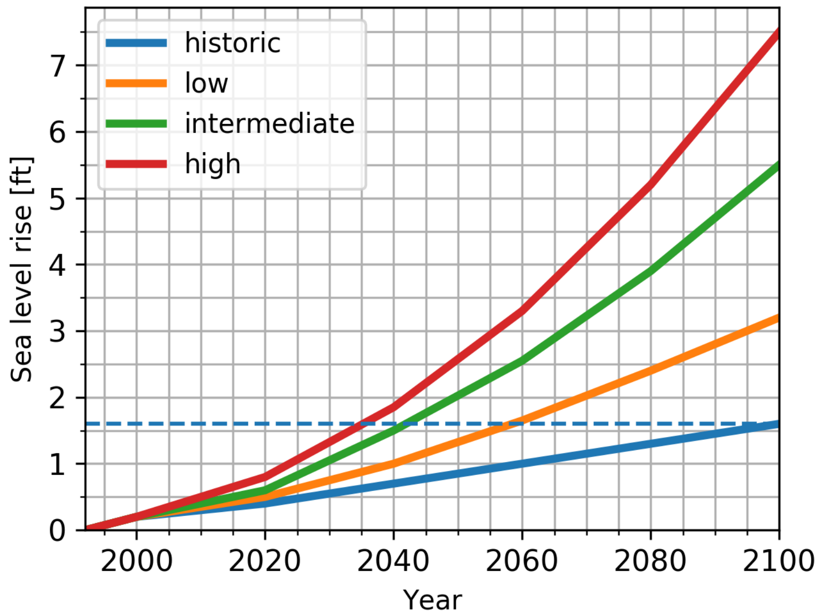

Figure 5.13: Region specific SLR scenarios for coastal Virginia. Adapted from (Mitchell et al., 2013)

line in Figure 5.13). When this threshold is crossed, I would expect to see street flooding from a typical daily tide cycle at high tide. Based on the SLR estimates from Mitchell et al. (2013), the 1.6 ft threshold could be crossed as early as the year 2035 for Norfolk in the high emission scenario and by 2060 for the low scenario.

Although the use of RTC for stormwater systems will increase with increasing average sea level in the coming decades, in current conditions storm surge and astronomical tides alone can elevate sea levels temporarily to the point where RTC would be useful. For example, since 2000, four storms, Hurricanes Isabel (2003), Irene (2011), and Sandy (2012), and a powerful Nor'easter (2009) each caused storm surge that reached over 3.93 ft (1.2 m). The astronomical tide alone has a 99% likelihood of reaching 2.85 ft (0.87 meters) each year in Norfolk (based on the 1983-2001 tidal epoch, see Figure 5.14). This is well above the 1.6 ft threshold, suggesting that Norfolk now experiences, and will continue to experience, flooding from the highest tide of each year. The results of this analysis suggest that with the addition of a tide gate, the introduction of active controls through MPC could significantly reduce flooding caused by flooding from similar storm surge and astronomical tide events.

5.4.2 General discussion

To make the results and inferences more complete, future work should be done to quantify uncertainty in the results. One source of uncertainty in this analysis include the model itself and its how well it represents the physical system. SWMM5 is a 1D model and therefore cannot simulate some aspects of coastal flooding. For

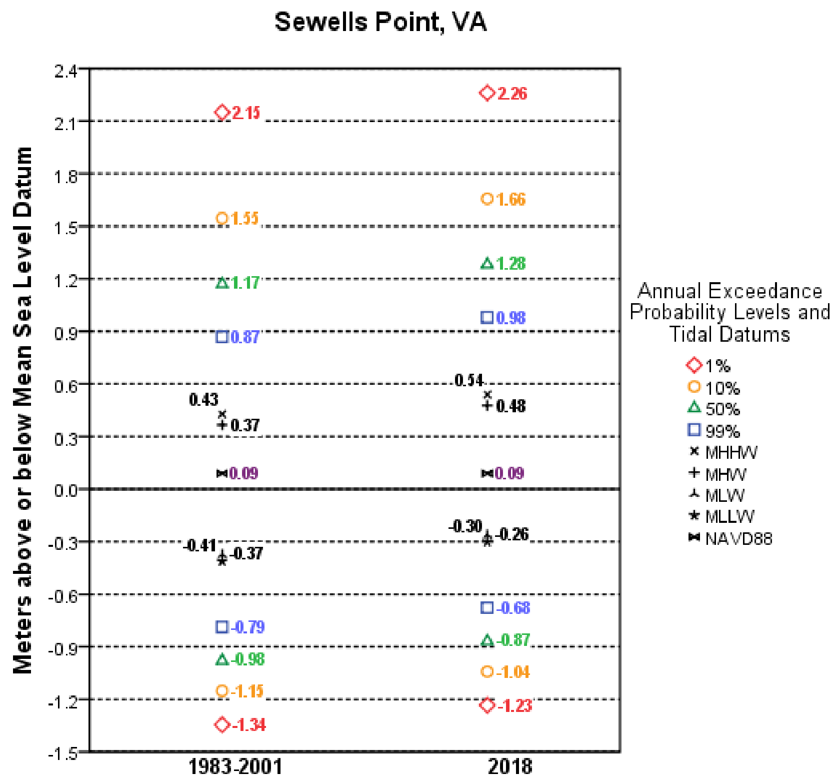

Figure 5.14: High and low exceedance probability levels for Sewell's Point tide gauge in Norfolk, VA USA (NOAA, 2017e)

example, SWMM5 only records a volume of water that is flooded from a given node, and cannot simulate the depth of flooding. Knowing the depth of flooding, and the reduction of depth by MPC, would be another way to assess the utility of MPC.

Another uncertainty in the results is in the optimality of the control strategies found in the MPC scenario by swmm_mpc. The SWMM5 model is used as a black-box model without the optimizer having any knowledge of the underlying mathematical structure of the governing equations. Because of this, the control strategies found by swmm_mpc cannot be guaranteed to be optimal. Future work could be done to analyze how close to optimal the strategies found through the swmm_mpc are.

In this analysis, I simulated three actuators. These were chosen in an ad hoc manner based on the behavior of the model and the constraints of the physical system. A methodology for systematically selecting which actuators should be placed and where would be useful extension of this work. Such a methodology could be especially useful as coastal cities seek for adaptation alternatives. The monetary cost for adding controls into the system could added to the cost function as an important practical consideration for municipalities. Optimizing the placement and selection of controls could be an addition to the swmm_mpc software.

Because the focus of this study was on increasing flooding with SLR, the rainfall magnitude and intensity were not changed. However, climate change, in addition to causing SLR, is anticipated to increase rainfall intensity (Berggren et al., 2012; Neumann et al., 2015). It would therefore be useful to, in addition to increasing sea level in the simulations, increase rainfall intensity and/or volume with time.A recent study has already been done by Dewberry to quantify increases in design storm rainfall volumes with climate change (Smirnov et al., 2018). By analyzing historic rainfall data and climate model the study recommends a 20% higher rainfall volumes for design storms in Virginia Beach. These projections could be used to adjust design storms and run additional scenarios using the methods of this study.

Barriers to adoption of RTC for stormwater systems by municipalities include costs and risk. A municipality must weigh the potential benefit (e.g., reduced flooding), with the costs and potential risks. The monetary costs include the costs of the equipment, installation, and maintenance. Given RTC systems' reliance on real-time communications, costs will include those associated with the information technology infrastructure including communication, data storage, and data access/display portals. There have been several RTC systems implemented for urban drainage systems or RTC for similar systems, therefore, estimating the costs should be relatively straightforward.

Less straightforward to estimate, and perhaps a larger barrier than monetary cost is that of potential risks of RTC systems in stormwater infrastructure. Although successfully deployed in some cities for combined sewer sytems, the use of RTC in separated sewer systems is still new. Schütze et al. (2004) provide some suggestions for overcoming institutional barriers for adopting RTC systems. Although primarily focused on combined sewer systems, many of the suggestions apply to separated sewer systesms as well. A key suggestion made by Schütze et al. (2004) in regards to safety and risk when implementing RTC is that in the worst-case scenario, the RTC system should perform as good or better than the system before the installation of the RTC. Additionally, Schütze et al. (2004) suggests that equipment failure should be considered inevitable and, therefore, redundancy should be built into the system. Finally, Schütze et al. (2004) emphasize the need to involve and educate municipality stormwater personnel from the beginning stages and potentially begin with an operator-in-the-loop approach until trust in the system is built.

5.5 Conclusions

In this research I assessed the utility of model predictive control (MPC) for reducing flooding in a coastal city with sea level rise (SLR). The study area was key neighborhood in Norfolk, VA, USA, a city particularly vulnerable to coastal flooding. A 2-year 12-hour design storm along with various tide scenarios were input into a SWMM5 model of the study area. This event was used because I was only considering non-extreme rainfall and tidal conditions (e.g., not hurricane-level events). To simulate a range of SLR scenarios, the base tide cycle was increased in 0.5 ft increments up to 3.5 ft.

The tidal increases were input for three scenarios, 1) a passive system (system as it currently behaves), 2) the passive system with the addition of a tide gate, and 3) the system with the addition of a tide gate and actuators controlled through MPC. Three actuators were simulated in the MPC scenario, a valve at the outlet of a pond, a pump to increase head, and an inflatable dam to utilize inline storage. The control strategies for the actuators were found using the swmm_mpc library with a genetic algorithm as the optimization method.

Flooding in the passive system increased dramatically after an increase in sea level of 1.6 ft. The addition of a tide gate greatly mitigated this increase in flooding. MPC further reduced overall flooding with an average effective percent reduction of 32%. The rate of the increase in flooding was significantly smaller with MPC than without. MPC also reduced maximum node flood volume by an average of 52% for sea level increases of 2.0 ft and above.

It is anticipated that SLR will make coastal cities more vulnerable to tidally- and rainfall-driven flooding. In addition to the installation of a tide gate, my results suggest that the use of actuators controlled by MPC could significantly reduce overall flood volumes and reduce flood severity at individual nodes in coastal cities. Transforming traditionally static, gravity-driven stormwater systems into dynamic, adaptive ones could reap large benefits in cumulative reduced flood volumes over the decades.

Acknowledgments

I thank the city of Norfolk for providing the SWMM5 model of the Hague area. I also acknowledge the support of awards 1735587 (Critical and Resilient Infrastructure Systems - CRISP) and 1737432 (Smart and Connected Communities - S&CC) from the United States National Science Foundation.

Chapter 6

Conclusion

Collectively, this dissertation advances the use of crowd-sourced data, data-driven techniques, and model predictive control for understanding, predicting, and mitigating urban coastal flooding. Such approaches can supplement the traditional approaches providing understanding and predictive power in coastal urban environments which are complex and can be difficult to model using physically-based models. Among the contributions are 1) a method for understanding the importance of rainfall observations close to flood-prone areas 2) a method for using crowd-sourced data and data-driven techniques for predicting flood severity 3) an open-source software package for simulating model predictive control for the EPA Stormwater Management Model and 4) an estimation of the current and future utility of active stormwater controls under different sea level rise scenarios.

This dissertation highlights the importance of nearby rain gauges for accurately estimating rainfall in a coastal city. In Chapter 2, data from the nearest quality controlled rain gauge and all closer rain gauges to flood-prone watersheds in Virginia Beach were removed to understand the effect of data from nearby rain gauges on rainfall estimation. Rainfall estimations changed on average by about 50% across all the watersheds at a 15-minute time step when the nearest station was excluded. For a single watershed, the highest average change in rainfall estimation was over 70% at a 15-minute time step with the largest difference in rainfall estimation of 24.5 mm at a 15-minute time step. This corresponds to the difference between a negligible design storm and a 10-year, 15-minute design storm. Results from this study also suggest that rain gauges added within 0.5 km can decrease prediction variance by 50-100% and a rain gauge 3.5 km from the watershed centroid will not decrease estimation variance appreciably.

The research in this dissertation also demonstrates that crowd-sourced data and data-driven models can be used for predicting urban, coastal street flood severity. In Chapter 3, two data-driven models, Poisson regression and Random Forest were trained to predict flood severity for a given set of environmental conditions (rainfall, tide levels, groundwater levels, and wind conditions) using quality-controlled, crowd-sourced street flooding reports as a proxy output variable. The data used for training and evaluating the models was from Norfolk, Virginia USA. The Random Forest model performed better overall compared to Poisson regression in the evaluation phase (root mean squared error of 3.87 compared to 6.71 flood reports, mean absolute error of 0.69 compared to 0.96 flood reports) with less variance (standard deviation of 6.00 compared to 18.42 flood reports). The most important variable in predicting model output in the Random Forest model was by far total cumulative rainfall followed by low tide and lower low tide.

As part of the research, I developed open-source software that can be used, built-from, and improved upon as a tool to assist decision-makers and researchers in finding effective control policies for urban drainage systems. Chapter 4, describes this software, swmm_mpc, which computes a control policy for controls within an urban drainage system model. The widely-used United States Environmental Protection Agency Stormwater Management Model Version 5 (SWMM5) was used to simulate the urban drainage system and the process model. An evolutionary algorithm was used to find an effective control policy at each time step. When tested using a simple SWMM5 model, the swmm_mpc software was able to produce control policies that met objectives including minimizing flooding and minimizing deviation from target water levels at certain nodes in the system.

The results from this dissertation suggest that, in addition to the installation of a tide gate, the use of actuators controlled by MPC could significantly reduce overall flood volumes and reduce flood severity at individual nodes in coastal cities as sea levels rise. This is addressed in Chapter 5. The study area was key neighborhood in Norfolk, VA, USA, a city particularly vulnerable to coastal flooding. Flooding in the passive system increased dramatically after an increase in sea level of 1.6 ft. The addition of a tide gate greatly mitigated this increase in flooding. MPC further reduced overall flooding with an average effective percent reduction of 32%. The rate of the increase in flooding was significantly smaller with MPC than without. MPC also reduced maximum node flood volume by an average of 52% for sea level increases of 2.0 ft and above.

Together, this dissertation provides methodologies for and demonstration of leveraging crowd-sourced data, data-driven techniques, and real-time control for improved adaptation in place to urban coastal flooding. Norfolk and Virginia Beach are experiencing the effects of sea level rise earlier than many coastal cities. Although demonstrated with these two cities, the methods developed and insight gained in this dissertation, with little modification, could be applied to and benefit other coastal cities who, in the coming decades, will likely experience similar sea level rise. As coastal flooding is predicted to increase in frequency, the cumulative benefit of better adaption to flooding could be very significant over time.

Bibliography

Adamovic, M., Branger, F., Braud, I., and Kralisch, S. (2016). Development of a data-driven semi-distributed hydrological model for regional scale catchments prone to Mediterranean flash floods. *Journal of Hydrology*, 541:173–189.

Adger, W. N. (2003). Social Capital, Collective Action, and Adaptation to Climate Change. *Economic Geography*, 79(4):387–404.

Afram, A. and Janabi-Sharifi, F. (2014). Theory and applications of HVAC control systems A review of model predictive control (MPC). *Building and Environment*, 72:343–355.

Alexander, L. V., Zhang, X., Peterson, T. C., Caesar, J., Gleason, B., Klein Tank, A. M. G., Haylock, M., Collins, D., Trewin, B., Rahimzadeh, F., Tagipour, A., Rupa Kumar, K., Revadekar, J., Griffiths, G., Vincent, L., Stephenson, D. B., Burn, J., Aguilar, E., Brunet, M., Taylor, M., New, M., Zhai, P., Rusticucci, M., and Vazquez-Aguirre, J. L. (2006). Global observed changes in daily climate extremes of temperature and precipitation. *Journal of Geophysical Research*, 111(D5):D05109.

Bartholmes, J. and Todini, E. (2005). Coupling meteorological and hydrological models for flood forecasting. *Hydrology and Earth System Sciences*, 9(4):333–346.

Basha, E. A., Ravela, S., and Rus, D. (2008). Model-based monitoring for early warning flood detection. In *Proceedings of the 6th ACM conference on Embedded network sensor systems - SenSys '08*, page 295, New York, New York, USA. ACM Press.

Bates, P. D., Dawson, R. J., Hall, J. W., Horritt, M. S., Nicholls, R. J., Wicks, J., and Mohamed Ahmed Ali Mohamed Hassan (2005). Simplified two-dimensional numerical modelling of coastal flooding and example applications. *Coastal Engineering*, 52(9):793–810.

Beguería, S. (2006). Validation and Evaluation of Predictive Models in Hazard Assessment and Risk Management. *Natural Hazards*, 37(3):315–329.

Bellon, A. and Austin, G. (1984). The accuracy of short-term radar rainfall forecasts. *Journal of Hydrology*, 70(1-4):35–49.

Berggren, K., Olofsson, M., Viklander, M., Svensson, G., and Gustafsson, A.-M. (2012). Hydraulic Impacts on Urban Drainage Systems due to Changes in Rainfall Caused by Climatic Change. *Journal of Hydrologic Engineering*, 17(1):92–98.

Berne, A., Delrieu, G., Creutin, J. D., and Obled, C. (2004). Temporal and spatial resolution of rainfall measurements required for urban hydrology. *Journal of Hydrology*, 299(3-4):166–179.

Bjerklie, D. M., Mullaney, J. R., Stone, J. R., Skinner, B. J., and Ramlow, M. A. (2012). Preliminary investigation of the effects of sea-level rise on groundwater levels in New Haven, Connecticut. *Open-File Report*.

Bonnin, G., Martin, D., Lin, B., Parzybok, T., Yekta, M., and Riley, D. (2018). Point Precipitation Frequency Estimates, Norfolk, Virginia, USA. https://hdsc.nws.noaa.gov/hdsc/pfds/pfds_printpage.html?lat=36.8661&lon=-76.2890&data=depth&units=english&series=pds.

Bonnin, G. M., Martin, D., Lin, B., Parzybok, T., Yekta, M., and Riley, D. (2006). NOAA Atlas 14: Precipitation-Frequency Atlas of the United States. http://hdsc.nws.noaa.gov/hdsc/pfds/pfds_printpage.html?lat=36.8472&lon=-75.9833&data=intensity&units=metric&series=pds.

Booker, D. J. and Woods, R. A. (2014). Comparing and combining physically-based and empirically-based approaches for estimating the hydrology of ungauged catchments. *Journal of Hydrology*, 508:227–239.

Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1):5–32.

Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). *Classification and Regression Trees.* Wadsworth, Belmont, CA.

Broder, J. M. (2009). Climate change seen as threat to US security. *The New York Times*, 9:1.

Burger, G., Sitzenfrei, R., Kleidorfer, M., and Rauch, W. (2014). Parallel flow routing in SWMM 5. *Environmental Modelling & Software*, 53:27–34.

Camacho, E. F. and Bordons, C. (2007). *Model predictive control*. Springer London Limited.

Cameron, A. C. and Trivedi, P. K. (1998). *Regression analysis of count data*, volume 53. Cambridge university press.

Caviedes-Voullième, D., García-Navarro, P., and Murillo, J. (2012). Influence of mesh structure on 2D full shallow water equations and SCS Curve Number simulation of rainfall/runoff events. *Journal of Hydrology*, 448-449:39–59.

Cembrano, G., Quevedo, J., Salamero, M., Puig, V., Figueras, J., and Mart, J. (2004). Optimal control of urban drainage systems. A case study. *Control Engineering Practice*, 12(1):1–9.

Church, J. A., Clark, P. U., Cazenave, A., J.M., G., Jevrejeva, S., Levermann, A., Merrifield, M., Milne, G., Nerem, R., Nunn, P., Payne, A., Pfeffer, W., Stammer, D., and Unnikrishnan, A. (2013). Sea level change. In Stocker, T., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P., editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Church, J. A., Gregory, J. M., Huybrechts, P., Kuhn, M., Lambeck, K., Nhuan, M. T., Qin, D., and Woodworth, P. L. (2001). Changes in sea level. In *in: JT Houghton, Y. Ding, DJ Griggs, M. Noguer, PJ Van der Linden, X. Dai, K. Maskell, and CA Johnson (eds.): Climate Change 2001: The Scientific Basis: Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel*, pages 639–694. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Ciach, G. J. and Krajewski, W. F. (2006). Analysis and modeling of spatial correlation structure in small-scale rainfall in Central Oklahoma. *Advances in Water Resources*, 29(10):1450–1463.

Collier, C. G. (2007). Flash flood forecasting: What are the limits of predictability? *Quarterly Journal of the Royal Meteorological Society*, 133(622):3–23.

Cooper, J. A. and Lemckert, C. (2012). Extreme sea-level rise and adaptation options for coastal resort cities: A qualitative assessment from the Gold Coast, Australia. *Ocean and Coastal Management*, 64:1–14.

Coxe, S., West, S. G., and Aiken, L. S. (2009). The Analysis of Count Data: A Gentle Introduction to Poisson Regression and Its Alternatives. *Journal of Personality Assessment*, 91(2):121–136.

Darsono, S. and Labadie, J. W. (2007). Neural-optimal control algorithm for real-time regulation of in-line storage in combined sewer systems. *Environmental Modelling & Software*, 22(9):1349–1361.

Del Re, L., Allgöwer, F., Glielmo, L., Guardiola, C., and Kolmanovsky, I. (2010). *Automotive model predictive control: models, methods and applications*, volume 402. Springer.

D'Unger, A., Land, K. C., McCall, P. L., and Nagin, D. S. (1998). How Many Latent Classes of Delinquent/Criminal Careers? Results from Mixed Poisson Regression Analyses. *American Journal of Sociology*, 103(6):1593–1630.

Emmanuel, I., Andrieu, H., Leblois, E., and Flahaut, B. (2012). Temporal and spatial variability of rainfall at the urban hydrological scale. *Journal of Hydrology*, 430-431:162–172.

Ercan, M. B. and Goodall, J. L. (2013). Estimating Watershed-Scale Precipitation by Combining Gauge- and Radar-Derived Observations. *Journal of Hydrologic Engineering*, 18(8):983–994.

Ezer, T. and Atkinson, L. P. (2014). Accelerated flooding along the U.S. East Coast: On the impact of sea-level rise, tides, storms, the Gulf Stream, and the North Atlantic Oscillations. *Earth's Future*, 2(8):362–382.

Fears, D. (2012). Built on sinking ground, Norfolk tries to hold back tide amid sea-level rise. *The Washington Post*. https://www.washingtonpost.com/national/health-science/built-on-sinking-ground-norfolk-tries-to-hold-back-tide-amid-sea-level-rise/2012/06/17/gJQADUsxjV_story.html?utm_term=.d93e580ed64e.

Fletcher, T., Andrieu, H., and Hamel, P. (2013). Understanding, management and modelling of urban hydrology and its consequences for receiving waters: A state of the art. *Advances in Water Resources*, 51:261–279.

Frome, E. L. and Checkoway, H. (1985). Use of Poisson Regression Models in Estimating Incidence Rates and Ratios. *American Journal of Epidemiology*, 121(2):309–323.

Gaitan, S., van de Giesen, N., and ten Veldhuis, J. (2016). Can urban pluvial flooding be predicted by open spatial data and weather data? *Environmental Modelling & Software*, 85:156–171.

Galarneau, T. J., Davis, C. A., and Shapiro, M. A. (2013). Intensification of Hurricane Sandy (2012) through Extratropical Warm Core Seclusion. *Monthly Weather Review*, 141(12):4296–4321.

Gallien, T. W., Sanders, B. F., and Flick, R. E. (2014). Urban coastal flood prediction: Integrating wave overtopping, flood defenses and drainage. *Coastal Engineering*, 91:18–28.

Gama, J. and Castillo, G. (2006). Learning with Local Drift Detection. In Li, X., Zaïane, O. R., and Li, Z., editors, *Advanced Data Mining and Applications: Second International Conference, ADMA 2006, Xi'an, China, August 14-16, 2006 Proceedings*, pages 42–55. Springer Berlin Heidelberg, Berlin, Heidelberg.

Gama, J., Žliobait, I., Bifet, A., Pechenizkiy, M., and Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4):1–37.

Gandomi, A. H., Yang, X.-s., Talatahari, S., and Alavi, A. H. (2013). *Metaheuristic Algorithms in Modeling and Optimization*. Elsevier.

Gardner, W., Mulvey, E. P., and Shaw, E. C. (1995). Regression Analyses of Counts and Rates: Poisson, Overdispersed Poisson, and Negative Binomial Models. *Psychological Bulletin*, 118(3):392–404.

Garofalo, G., Giordano, A., Piro, P., Spezzano, G., and Vinci, A. (2017). A distributed real-time approach for mitigating CSO and flooding in urban drainage systems. *Journal of Network and Computer Applications*, 78:30–42.

Gelormino, M. S. and Ricker, N. L. (1994). Model-predictive control of a combined sewer system. *International Journal of Control*, 59(3):793–816.

Godschalk, D. R. (2003). Urban Hazard Mitigation: Creating Resilient Cities. *Natural Hazards Review*, 4(3):136–143.

Goovaerts, P. (2000). Geostatistical approaches for incorporating elevation into the spatial interpolation of rainfall. *Journal of Hydrology*, 228(1-2):113–129.

Hall, M. (1984). *Urban hydrology*. Elsevier. Applied Science Publishers, New York.

Hallegatte, S. (2009). Strategies to adapt to an uncertain climate change. *Global Environmental Change*, 19(2):240–247.

Hayward, D. and Clarke, R. T. (1996). Relationship between rainfall, altitude and distance from the sea in the Freetown Peninsula, Sierra Leone. *Hydrological Sciences Journal*, 41(3):377–384.

He, H. and Garcia, E. A. (2009). Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284.

Heusch, S. and Ostrowski, M. (2011). Model Predictive Control with SWMM. *Journal of Water Management Modeling*, 19:237–247.

Hill, D., Kerkez, B., Rasekh, A., Ostfeld, A., Minsker, B., and Banks, M. K. (2014). Sensing and Cyberinfrastructure for Smarter Water Management: The Promise and Challenge of Ubiquity. *Journal of Water Resources Planning and Management*, 140(7):01814002.

Hill, D. J. (2015). Assimilation of weather radar and binary ubiquitous sensor measurements for quantitative precipitation estimation. *Journal of Hydroinformatics*, 17(4):598.

Homer, C. G., Dewitz, J. A., Yang, L., Jin, S., Danielson, P., Xian, G., Coulston, J., Herold, N. D., Wickham, J., and Megown, K. (2015). Completion of the 2011 National Land Cover Database for the conterminous United States-Representing a decade of land cover change information. *Photogramm. Eng. Remote Sens*, 81(5):345—-354.

Hong, W.-C. and Pai, P.-F. (2007). Potential assessment of the support vector regression technique in rainfall forecasting. *Water Resources Management*, 21(2):495–513.

Hoover, D. J., Odigie, K. O., Swarzenski, P. W., and Barnard, P. (2016). Sea-level rise and coastal groundwater inundation and shoaling at select sites in California, USA. *Journal of Hydrology: Regional Studies.*

Horsburgh, J. S., Tarboton, D. G., Piasecki, M., Maidment, D. R., Zaslavsky, I., Valentine, D., and Whitenack, T. (2009). An integrated system for publishing environmental observations data. *Environmental Modelling & Software*, 24(8):879–888.

Huber, W. C., Rossman, L. A., and Dickinson, R. E. (2005). EPA storm water management model, SWMM5. *Watershed Modeling, CRC Press, Boca Raton, FL*, pages 339–361.

Hunter, N., Bates, P., Neelz, S., Pender, G., Villanueva, I., Wright, N., Liang, D., Falconer, R., Lin, B., Waller, S., Crossley, A., and Mason, D. (2008). Benchmarking 2D hydraulic models for urban flood simulations. *Proceedings of the institution of civil engineers: water management*, 161(1):13—-30.

Ireland, G., Volpi, M., and Petropoulos, G. (2015). Examining the Capability of Supervised Machine Learning Classifiers in Extracting Flooded Areas from Landsat TM Imagery: A Case Study from a Mediterranean Flood. *Remote Sensing*, 7(3):3372–3399.

Jensen, N. and Pedersen, L. (2005). Spatial variability of rainfall: Variations within a single radar pixel. *Atmospheric Research*, 77(1):269–277.

Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research*, 4:237–285.

Kates, R. W., Colten, C. E., Laska, S., and Leatherman, S. P. (2006). Reconstruction of New Orleans after Hurricane Katrina: a research perspective. *Proceedings of the National Academy of Sciences of the United States of America*, 103(40):14653–60.

Kebaili Bargaoui, Z. and Chebbi, A. (2009). Comparison of two kriging interpolation methods applied to spatiotemporal rainfall. *Journal of Hydrology*, 365(1):56–73.

Kerkez, B., Gruden, C., Lewis, M., Montestruque, L., Quigley, M., Wong, B., Bedig, A., Kertesz, R., Braun, T., Cadwalader, O., Poresky, A., and Pak, C. (2016). Smarter Stormwater Systems. *Environmental Science & Technology*, 50(14):7267–7273.

Kim, I. Y. and de Weck, O. L. (2005). Adaptive weighted-sum method for bi-objective optimization: Pareto front generation. *Structural and multidisciplinary optimization*, 29(2):149–158.

Kleinosky, L. R., Yarnal, B., and Fisher, A. (2006). Vulnerability of Hampton Roads, Virginia to Storm-Surge Flooding and Sea-Level Rise. *Natural Hazards*, 40(1):43–70.

Krajewski, W. and Smith, J. (2002). Radar hydrology: rainfall estimation. *Advances in Water Resources*, 25(8):1387–1394.

Krajewski, W. F., Ciach, G. J., and Habib, E. (2003). An analysis of small-scale rainfall variability in different climatic regimes. *Hydrological Sciences Journal*, 48(2):151–162.

Kuhn, M., Wing, J., Weston, S., Williams, A., Keefer, C., Engelhardt, A., Cooper, T., Mayer, Z., and Kenkel, B. (2016). caret: Classification and Regression Training.

Lambert, D. (1992). Zero-Inflated Poisson Regression, with an Application to Defects in Manufacturing. *Technometrics*, 34(1):1.

Langson, W., Chryssochoos, I., Raković, S., and Mayne, D. (2004). Robust model predictive control using tubes. *Automatica*, 40(1):125–133.

Le Coz, J., Patalano, A., Collins, D., Guillen, N. F., Garcia, C. M., Smart, G. M., Bind, J., Chiaverini, A., Le Boursicaud, R., Dramais, G., and Braud, I. (2016). Crowdsourced data for flood hydrology: Feedback from recent citizen science projects in Argentina, France and New Zealand. *Journal of Hydrology*, 541:766–777.

Leandro, J., Chen, A. S., Djordjević, S., and Savić, D. A. (2009). Comparison of 1D/1D and 1D/2D Coupled (Sewer/Surface) Hydraulic Models for Urban Flood Simulation. *Journal of Hydraulic Engineering*, 135(6):495–504.

Liaw, A. and Wiener, M. (2002). Classification and Regression by randomForest. *R News*, 2(3):18–22.

Loh, W.-Y. (2011). Classification and regression trees. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1):14–23.

Loos, M. and Elsenbeer, H. (2011). Topographic controls on overland flow generation in a forest An ensemble tree approach. *Journal of Hydrology*, 409(1-2):94–103.

Maier, H., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L., Cunha, M., Dandy, G., Gibbs, M., Keedwell, E., Marchi, A., Ostfeld, A., Savic, D., Solomatine, D., Vrugt, J., Zecchin, A., Minsker, B., Barbour, E., Kuczera, G., Pasha, F., Castelletti, A., Giuliani, M., and Reed, P. (2014). Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions. *Environmental Modelling & Software*, 62:271–299.

Marfai, M. A., King, L., Sartohadi, J., Sudrajat, S., Budiani, S. R., and Yulianto, F. (2008). The impact of tidal flooding on a coastal community in Semarang, Indonesia. *Environmentalist*, 28(3):237–248.

Mark, O., Weesakul, S., Apirumanekul, C., Aroonnet, S. B., and Djordjevic, S. (2004a). Potential and limitations of 1D modelling of urban flooding. *Journal of Hydrology*, 299(3-4):284–299.

Mark, O., Weesakul, S., Apirumanekul, C., Aroonnet, S. B., and Djordjević, S. (2004b). Potential and limitations of 1D modelling of urban flooding. *Journal of Hydrology*, 299(3-4):284–299.

Mayne, D., Seron, M., and Raković, S. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224.

Meneses, E., Gaussens, M., Jakobsen, C., Mikkelsen, P., Grum, M., and Vezzaro, L. (2018). Coordinating Rule-Based and System-Wide Model Predictive Control Strategies to Reduce Storage Expansion of Combined Urban Drainage Systems: The Case Study of Lundtofte, Denmark. *Water*, 10(1):76.

Mignot, E., Paquier, A., and Haider, S. (2006). Modeling floods in a dense urban area using 2D shallow water equations. *Journal of Hydrology*, 327(1-2):186–199.

Mitchell, M., Hershner, C., Julie, H., Schatt, D., Mason, P., and Eggington, E. (2013). Recurrent Flooding Study for Tidewater Virginia. *Virginia Institute of Marine Science, Center for Coastal Resources Management*, (January).

Mockus, V. (2012). Storm Rainfall Depth and Distribution. In *National Engineering Handbook*, chapter 4. Natural Resources Conservation Service.

Moftakhari, H. R., AghaKouchak, A., Sanders, B. F., Feldman, D. L., Sweet, W., Matthew, R. A., and Luke, A. (2015). Increased nuisance flooding along the coasts of the United States due to sea level rise: Past and future. *Geophysical Research Letters*, 42(22):9846–9852.

Morrison, J. E. and Smith, J. A. (2002). Stochastic modeling of flood peaks using the generalized extreme value distribution. *Water Resources Research*, 38(12):41–1–41–12.

Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.

Mutanga, O., Adam, E., and Cho, M. A. (2012). High density biomass estimation for wetland vegetation using WorldView-2 imagery and random forest regression algorithm. *International Journal of Applied Earth Observation and Geoinformation*, 18:399–406.

Naghibi, S. A., Moghaddam, D. D., Kalantar, B., Pradhan, B., and Kisi, O. (2017). A comparative assessment of GIS-based data mining models and a novel ensemble model in groundwater well potential mapping. *Journal of Hydrology*, 548:471–483.

Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The computer journal*, 7(4):308–313.

Nesbitt, S. W. and Anders, A. M. (2009). Very high resolution precipitation climatologies from the Tropical Rainfall Measuring Mission precipitation radar. *Geophysical Research Letters*, 36(15).

Neumann, J. E., Price, J., Chinowsky, P., Wright, L., Ludwig, L., Streeter, R., Jones, R., Smith, J. B., Perkins, W., Jantarasami, L., and Martinich, J. (2015). Climate change risks to US infrastructure: impacts on roads, bridges, coastal development, and urban drainage. *Climatic Change*, 131(1):97–109.

Nicholls, R. J. and Cazenave, A. (2010). Sea-level rise and its impact on coastal zones. *Science (New York, N.Y.)*, 328(5985):1517–20.

NOAA (2017a). Daily Summaries Station Details: NORFOLK INTERNATIONAL AIRPORT, VA US, GHCND:USW00013737 — Climate Data Online (CDO) — National Climatic Data Center (NCDC). https://www.ncdc.noaa.gov/cdo-web/datasets/GHCND/stations/GHCND:USW00013737/detail.

NOAA (2017b). Daily Summaries Station Details: NORFOLK NAS, VA US, GHCND:USW00013750 — Climate Data Online (CDO) — National Climatic Data Center (NCDC). https://www.ncdc.noaa.gov/cdo-web/datasets/GHCND/stations/GHCND:USW00013750/detail.

NOAA (2017c). Money Point - Station Home Page - NOAA Tides & Currents. http://tidesandcurrents.noaa.gov/stationhome.html?id=8518750#info.

NOAA (2017d). NOAA Tide Predictions Users Guide. https://tidesandcurrents.noaa.gov/PageHelp.html.

NOAA (2017e). Sewells Point - Station Home Page - NOAA Tides & Currents. https://tidesandcurrents.noaa.gov/stationhome.html?id=8638610.

NOAA (2017f). What percentage of the American population lives near the coast? https://oceanservice.noaa.gov/facts/population.html.

O'Gorman, P. A. and Schneider, T. (2009). The physical basis for increases in precipitation extremes in simulations of 21st-century climate change. *Proceedings of the National Academy of Sciences*, 106(35):14773–14777.

Osgood, D. W. (2000). Poisson-Based Regression Analysis of Aggregate Crime Rates. *Journal of Quantitative Criminology*, 16(1):21–43.

Overeem, A., Leijnse, H., and Uijlenhoet, R. (2013). Country-wide rainfall maps from cellular communication networks. *Proceedings of the National Academy of Sciences of the United States of America*, 110(8):2741–5.

Pedersen, L., Jensen, N. E., Christensen, L. E., and Madsen, H. (2010). Quantification of the spatial variability of rainfall based on a dense network of rain gauges. *Atmospheric Research*, 95(4):441–454.

Pleau, M., Colas, H., Lavallée, P., Pelletier, G., and Bonin, R. (2005). Global optimal real-time control of the Quebec urban drainage system.

Povak, N. A., Hessburg, P. F., McDonnell, T. C., Reynolds, K. M., Sullivan, T. J., Salter, R. B., and Cosby, B. J. (2014). Machine learning and linear regression models to predict catchment-level base cation weathering rates across the southern Appalachian Mountain region, USA. *Water Resources Research*, 50(4):2798–2814.

Puig, V., Cembrano, G., Romera, J., Quevedo, J., Aznar, B., Ramón, G., and Cabot, J. (2009). Predictive optimal control of sewer networks using CORAL tool: application to Riera Blanca catchment in Barcelona. *Water Science & Technology*, 60(4):869.

Qin, S. and Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733–764.

R Core Team (2017). R: A Language and Environment for Statistical Computing. https://www.r-project.org/.

Ray, R. D. and Foster, G. (2016). Future nuisance flooding at Boston caused by astronomical tides alone. *Earth's Future*.

Renard, D., Bex, N., Desassis, N., Beucher, H., Ors, F., and Laporte, F. (2015). RGeostats: The Geostatistical package.

Rosenzweig, C., Solecki, W. D., Blake, R., Bowman, M., Faris, C., Gornitz, V., Horton, R., Jacob, K., LeBlanc, A., Leichenko, R., Linkin, M., Major, D., O'Grady, M., Patrick, L., Sussman, E., Yohe, G., and Zimmerman, R. (2011). Developing coastal adaptation to climate change in the New York City infrastructure-shed: process, approach, tools, and strategies. *Climatic Change*, 106(1):93–127.

Rotzoll, K. and Fletcher, C. H. (2012). Assessment of groundwater inundation as a consequence of sea-level rise. *Nature Climate Change*, 3(5):477–481.

Sadler, J., Goodall, J., Morsy, M., and Spencer, K. (2018a). Modeling urban coastal flood severity from crowd-sourced flood reports using Poisson regression and Random Forest. *Journal of Hydrology*, 559:43–55.

Sadler, J. M., Goodall, J. L., Morsy, M. M., and Spencer, K. (2018b). Modeling Urban Coastal Flood Severity from Crowd-Sourced Flood Reports Using Poisson Regression and Random Forest. IN REVISION. *Journal of Hydrology*.

Sadler, J. M., Haselden, N., Mellon, K., Hackel, A., Son, V., Mayfield, J., Blase, A., and Goodall, J. L. (2017). Impact of Sea-Level Rise on Roadway Flooding in the Hampton Roads Region, Virginia. *Journal of Infrastructure Systems*, 23(4):05017006.

Sahoo, S., Russo, T. A., Elliott, J., and Foster, I. (2017). Machine learning algorithms for modeling groundwater level changes in agricultural regions of the U.S. *Water Resources Research*, 53(5):3878–3895.

Schütze, M., Campisano, A., Colas, H., Schilling, W., and Vanrolleghem, P. A. (2004). Real time control of urban wastewater systemswhere do we stand today? *Journal of Hydrology*, 299(3-4):335–348.

Seo, D.-J. (1998). Real-time estimation of rainfall fields using radar rainfall and rain gage data. *Journal of Hydrology*, 208(1-2):37–52.

Serinaldi, F. (2008). Analysis of inter-gauge dependence by Kendall's τK, upper tail dependence coefficient, and 2-copulas with application to rainfall fields. *Stochastic Environmental Research and Risk Assessment*, 22(6):671–688.

Smirnov, D., Giovannettone, J., Lawler, S., Sreetharan, M., Plummer, J., and Workman, B. (2018). Analysis of historical and future heavy precipitation: City of virginia beach, virginia. Technical report.

Smith, J. A., Baeck, M. L., Meierdiercks, K. L., Miller, A. J., and Krajewski, W. F. (2007). Radar rainfall estimation for flash flood forecasting in small urban watersheds. *Advances in Water Resources*, 30(10):2087–2097.

Smith, J. A. and Krajewski, W. F. (1993). A modeling study of rainfall ratereflectivity relationships. *Water Resources Research*, 29(8):2505–2514.

Smith, R. A., Bates, P. D., and Hayes, C. (2011). Evaluation of a coastal flood inundation model using hard and soft data. *Environmental Modelling & Software*, 30:35–46.

Smith, T. L., Benjamin, S. G., Brown, J. M., Weygandt, S., Smirnova, T., and Schwartz, B. (2008). Convection Forecasts from the hourly Updated, 3-KM High Resolution Rapid Refresh (HRRR) Model. In *Preprints, 24th Conference on Severe Local Storms*, Savannah, GA.

Solomatine, D. P. and Ostfeld, A. (2008). Data-driven modelling: some past experiences and new approaches. *Journal of Hydroinformatics*, 10(1):3.

Solomatine, D. P. and Xue, Y. (2004). M5 Model Trees and Neural Networks: Application to Flood Forecasting in the Upper Reach of the Huai River in China. *Journal of Hydrologic Engineering*, 9(6):491–501.

Suarez, P., Anderson, W., Mahal, V., and Lakshmanan, T. (2005). Impacts of flooding and climate change on urban transportation: A systemwide performance assessment of the Boston Metro Area. *Transportation Research Part D: Transport and Environment*, 10(3):231–244.

Sun, X., Mein, R., Keenan, T., and Elliott, J. (2000). Flood estimation using radar and raingauge data. *Journal of Hydrology*, 239(1-4):4–18.

Svetnik, V., Liaw, A., Tong, C., Culberson, J. C., Sheridan, R. P., and Feuston, B. P. (2003). Random forest: a classification and regression tool for compound classification and QSAR modeling. *Journal of chemical information and computer sciences*, 43(6):1947–1958.

Sweet, W., Park, J., Marra, J., Zervas, C., and Gill, S. (2014). Sea Level Rise and Nuisance Flood Frequency Changes around the United States. *NOAA Technical Report NOS CO-OPS 073*, (June):58.

Sweet, W. V. and Park, J. (2014). From the extreme to the mean: Acceleration and tipping points of coastal inundation from sea level rise. *Earth's Future*, 2(12):579–600.

Tao, Y., Yang, T., Faridzad, M., Jiang, L., He, X., and Zhang, X. (2017). Non-stationary bias correction of monthly CMIP5 temperature projections over China using a residual-based bagging tree model.

Tayfur, G., Kavvas, M. L., Govindaraju, R. S., and Storm, D. E. (1993). Applicability of St. Venant Equations for TwoDimensional Overland Flows over Rough Infiltrating Surfaces. *Journal of Hydraulic Engineering*, 119(1):51–63.

Tehrany, M. S., Pradhan, B., and Jebur, M. N. (2013). Spatial prediction of flood susceptible areas using rule based decision tree (DT) and a novel ensemble bivariate and multivariate statistical models in GIS. *Journal of Hydrology*, 504:69–79.

The Port of Virginia (2016). Hampton Roads Harbor, 2016 Trade Overview. http://www.portofvirginia.com/pdfs/about/POV2016_Trade_Overview.pdf.

Tien Bui, D., Pradhan, B., Nampak, H., Bui, Q.-T., Tran, Q.-A., and Nguyen, Q.-P. (2016). Hybrid artificial intelligence approach based on neural fuzzy inference model and metaheuristic optimization for flood susceptibilitgy modeling in a high-frequency tropical cyclone area using GIS. *Journal of Hydrology*, 540:317–330.

Tiernan, B. (2017). 4 things we learned from tracing the king tide across Hampton Roads. https://pilotonline.com/news/local/environment/things-we-learned-from-tracing-the-king-tide-across-hampton/article_81e82d0d-b6d6-53c7-a1da-a94b51305e07.html.

Titus, J. G., Kuo, C. Y., Gibbs, M. J., LaRoche, T. B., Webb, M. K., and Waddell, J. O. (1987). Greenhouse Effect, Sea Level Rise, and Coastal Drainage Systems. *Journal of Water Resources Planning and Management*, 113(2):216–227.

Turner, R. E. (2004). Coastal wetland subsidence arising from local hydrologic manipulations. *Estuaries*, 27(2):265–272.

U.S. Department of Commerce (2012). Virginia: 2010, Population and Housing Unit Counts. (July).

Valverde Ramírez, M. C., de Campos Velho, H. F., and Ferreira, N. J. (2005). Artificial neural network technique for rainfall forecasting applied to the São Paulo region. *Journal of Hydrology*, 301(1-4):146–162.

Velasco-Forero, C. A., Sempere-Torres, D., Cassiraga, E. F., and Jaime Gómez-Hernández, J. (2009). A non-parametric automatic blending methodology to estimate rainfall fields from rain gauge and radar data. *Advances in Water Resources*, 32(7):986–1002.

Vermeer, M. and Rahmstorf, S. (2009). Global sea level linked to global temperature. *Proceedings of the National Academy of Sciences of the United States of America*, 106(51):21527–21532.

Viglione, A., Di Baldassarre, G., Brandimarte, L., Kuil, L., Carr, G., Salinas, J. L., Scolobig, A., and Blöschl, G. (2014). Insights from socio-hydrology modelling on dealing with flood risk Roles of collective memory, risk-taking attitude and trust. *Journal of Hydrology*, 518:71–82.

Vrabie, D., Pastravanu, O., Abu-Khalaf, M., and Lewis, F. (2009). Adaptive optimal control for continuous-time linear systems based on policy iteration. *Automatica*, 45(2):477–484.

Wang, W.-C., Chau, K.-W., Cheng, C.-T., and Qiu, L. (2009). A comparison of performance of several artificial intelligence methods for forecasting monthly discharge time series. *Journal of Hydrology*, 374(3):294–306.

Wang, Z., Lai, C., Chen, X., Yang, B., Zhao, S., and Bai, X. (2015). Flood hazard risk assessment model based on random forest. *Journal of Hydrology*, 527:1130–1141.

Wong, B. P. and Kerkez, B. (2018). RealTime Control of Urban Headwater Catchments Through Linear Feedback: Performance, Analysis, and Site Selection. *Water Resources Research*, 54(10):7309–7330.

Wu, W., Kitzmiller, D., Wu, S., Wu, W., Kitzmiller, D., and Wu, S. (2012). Evaluation of Radar Precipitation Estimates from the National Mosaic and Multisensor Quantitative Precipitation Estimation System and the WSR-88D Precipitation Processing System over the Conterminous United States. *Journal of Hydrometeorology*, 13(3):1080–1093.

Xu, T., Valocchi, A. J., Ye, M., and Liang, F. (2017). Quantifying model structural error: Efficient Bayesian calibration of a regional groundwater flow model using surrogates and a data-driven error model. *Water Resources Research*.

Yang, T., Asanjan, A. A., Faridzad, M., Hayatbini, N., Gao, X., and Sorooshian, S. (2017a). An enhanced artificial neural network with a shuffled complex evolutionary global optimization with principal component analysis. *Information Sciences*, 418-419:302–316.

Yang, T., Asanjan, A. A., Welles, E., Gao, X., Sorooshian, S., and Liu, X. (2017b). Developing reservoir monthly inflow forecasts using artificial intelligence and climate phenomenon information. *Water Resources Research*, 53(4):2786–2812.

Yang, T., Gao, X., Sorooshian, S., and Li, X. (2016). Simulating California reservoir operation using the classification and regression-tree algorithm combined with a shuffled cross-validation scheme. *Water Resources Research*, 52(3):1626–1651.

Zevenbergen, C., Veerbeek, W., Gersonius, B., and Van Herk, S. (2008). Challenges in urban flood management: travelling across spatial and temporal scales. *Journal of Flood Risk Management*, 1(2):81–88.

Zimmer, A., Schmidt, A., Ostfeld, A., and Minsker, B. (2015). Evolutionary algorithm enhancement for model predictive control and real-time decision support. *Environmental Modelling & Software*, 69:330–341.

Zimmer, A., Schmidt, A., Ostfeld, A., and Minsker, B. (2018). Reducing Combined Sewer Overflows through Model Predictive Control and Capital Investment. *Journal of Water Resources Planning and Management*, 144(2):04017091.

Zinevich, A., Alpert, P., and Messer, H. (2008). Estimation of rainfall fields using commercial microwave communication networks of variable density. *Advances in Water Resources*, 31(11):1470–1480.