Title: Rapid prediction of peak storm surge from tropical cyclone track time series using machine learning

Authors: Jun-Whan Lee (jwlee89@vt.edu)[a], Jennifer L. Irish[a,b], Michelle T. Bensi[c] and Douglas C. Marcy[d]

a Department of Civil and Environmental Engineering, Virginia Tech, 750 Drillfield Dr, Blacksburg, VA, 24061, USA

b Center for Coastal Studies, Virginia Tech, 926 W Campus Dr, Blacksburg, VA 24061

c Department of Civil and Environmental Engineering, University of Maryland, College Park, MD 20742, USA

d NOAA Oce for Coastal Management, 2234 S. Hobson Avenue, Charleston, SC 29405, USA

Twitter: @junwhanlee89

Rapid prediction of peak storm surge from tropical cyclone track time series using machine learning

Jun-Whan Lee[a,*], Jennifer L. Irish[a,b], Michelle T. Bensi[c] and Douglas C. Marcy[d]

[a]Department of Civil and Environmental Engineering, Virginia Tech, 750 Drillfield Dr, Blacksburg, VA, 24061, USA

[b]Center for Coastal Studies, Virginia Tech, 926 W Campus Dr, Blacksburg, VA 24061, USA

[c]Department of Civil and Environmental Engineering, University of Maryland, College Park, MD 20742, USA

[d]NOAA Office for Coastal Management, 2234 S. Hobson Avenue, Charleston, SC 29405, USA

ARTICLE INFO

ABSTRACT

Rapid and accurate prediction of peak storm surges across an extensive coastal region is necessary to inform assessments used to design the systems that protect coastal communities' life and property. Significant advances in high-fidelity, physics-based numerical models have been made in recent years, but use of these models for probabilistic forecasting and probabilistic hazard assessment is computationally intensive. Several surrogate modeling approaches based on existing databases of high-fidelity synthetic storm surge simulations have been recently suggested to reduce computational burden without substantial loss of accuracy. In these previous studies, however, the surrogate modeling approaches relied on a tropical cyclone condition at one moment (usually at or near landfall), which is not always most correlated with the peak storm surge. In this study, a new one-dimensional convolutional neural network model combined with principal component analysis and a k-means clustering (C1PKNet) is presented that can rapidly predict peak storm surge across an extensive coastal region from time-series of tropical cyclone conditions, namely the storm track. The C1PKNet model was trained and cross-validated for the Chesapeake Bay area of the United States using existing database of 1031 high-fidelity storm surge simulations, including both landfalling and bypassing storms. Moreover, the performance of the C1PKNet model was evaluated based on observations from three historical hurricanes (Hurricane Isabel in 2003, Hurricane Irene in 2011, and Hurricane Sandy in 2012). The results indicate that the C1PKNet model is computationally efficient and can predict peak storm surges from realistic tropical cyclone track time-series. We believe that this new surrogate model can enhance coastal resilience by providing rapid storm surge predictions.

1. Introduction

Storm surge generated by tropical cyclones (TC) is one of the most devastating threats to coastal communities. In the United States (US) alone, storm surge events have caused hundreds of billions of dollars in economic damage and have been responsible for about half of the fatalities from Atlantic TCs (1963-2012) (Rappaport, 2014). Historically, the US's costliest natural disaster was Hurricane Katrina (2005), where the storm surge exceeded 7.5 m above sea level, causing at least $170 billion of property damage (Robertson et al., 2007; Blake et al., 2011; Smith, 2020). Furthermore, the single deadliest natural disaster in the US was the unnamed Galveston hurricane (1900) that caused at least 8,000 deaths, most of which were storm-surge related (Blake et al., 2011). To mitigate the damages that can be caused by storm surge, several studies have emphasized the need for a rapid and accurate storm surge model that can support not only probabilistic storm surge forecasting, where surge predictions for a large number of storms are required in order to characterize uncertainty associated with uncertainty in the storm track (Glahn et al., 2009), but also robust probabilistic storm surge hazard assessment, where surge predictions for a large number of storms are also required to both accurately characterize mean hazard statistics and to account for uncertainty in the hazard estimation (Resio et al., 2007).

Physics-based numerical models, such as SLOSH (Sea, Lake, and Overland Surges from Hurricanes, Jelesnianski, 1992) and ADCIRC (ADvanced CIRCulation, Luettich et al., 1992), are widely used to predict storm surges. Over the last several decades, significant advances in physics-based model frameworks allow highly accurate simulation of storm surge by both accounting for dominant processes–e.g., wind surge, wave setup, nonlinear advection–and by employing

---

*Corresponding author
✉ jwlee89@vt.edu (J. Lee)
ORCID(s): 0000-0002-6722-6843 (J. Lee); 0000-0002-2429-5953 (J.L. Irish)

Table 1
Summary of representative surrogate modeling studies to predict peak storm surges from tropical cyclone information.

Study	Surrogate model type	Study area (Basin)	TC track time-series	Multi-Output	Overland	Bypassing storm	Validation
Irish et al. (2009)	SRFs	Texas (Northern Gulf of Mexico)	No	No	No	No	Synthetic
Al Kajbaf and Bensi (2020)	MLP, SVR, kriging	New Jersey (North Atlantic)	No	No	No	Yes	Synthetic
Hashemi et al. (2016)	MLP, SVR	Rhode Island (North Atlantic)	No	Yes[1]	No	No	Synthetic+Historical
Jia and Taflanidis (2013)	kriging+PCA	Hawaii (Central Pacific)	No	Yes[2]	No	N/A	Synthetic
Taflanidis et al. (2013b)	MLS	Hawaii (Central Pacific)	No	Yes[2]	Yes∗	N/A	Synthetic
Taflanidis et al. (2013a)	MLS	Hawaii (Central Pacific)	No	Yes[2]	Yes∗∗	N/A	Synthetic
Jia et al. (2016)	kriging+PCA	Louisiana (Northern Gulf of Mexico)	No	Yes[2]	Yes∗∗	No	Synthetic
Bass and Bedient (2018)	kriging+PCA, MLP+PCA	Texas (Northern Gulf of Mexico)	No	Yes[2]	Yes∗∗	No	Synthetic
Kyprioti et al. (2021)	kriging+PCA+KMC	New Jersey, New York (North Atlantic)	No	Yes[2]	Yes∗∗∗	No	Synthetic
This study	CNN+PCA+KMC	Virginia, Maryland (North Atlantic)	Yes	Yes[2]	Yes∗∗∗	Yes	Synthetic+Historical

SRFs = Surge response functions; MLP = Multilayer perceptron; SVR = Support vector regression; PCA = Principal component analysis; MLS = Moving least squares response surface; KMC = K-means clustering; CNN = Convolutional neural network; 1 = Two locations; 2 = >1,000 locations; ∗ = One-dimensional Boussinesq model; ∗∗ = Taflanidis et al. (2013a)'s method; ∗∗∗ = k-nearest neighbor interpolation; N/A = Not applicable.

high grid resolution. The problem, however, is that high-fidelity applications of physics-based numerical models are computationally intensive. This leads to challenges when using high-fidelity models to generate the large number of model runs required for probabilistic forecasting and probabilistic hazard assessment, even with parallelization techniques. For this reason, the National Hurricane Center's operational probabilistic forecast system (P-Surge) relies on a large number rapid, low-fidelity simulations that use a coarse-resolution grid within the more simplified SLOSH storm surge model that neglects key physical components such as nonlinear advection (Kerr et al., 2013). Moreover, due to the computational burden of the high-fidelity physics-based numerical modeling approach, most probabilistic storm surge hazard assessments have conventionally relied on either a limited number of high-fidelity storm surge simulations (e.g., Krien et al., 2015) or a large number of low-fidelity storm surge simulations (e.g., Lin et al., 2010); Resio and Irish (2015) and Resio et al. (2017) highlight shortcomings of both approaches.

Over the last decade, data-driven surrogate modeling based on high-fidelity synthetic storm simulation databases (e.g., IPET, 2006; Kennedy et al., 2012; Cialone et al., 2015) has gained increased attention. Namely, several studies suggest that the surrogate model approach can efficiently generate storm surge predictions without substantial loss of accuracy while drastically reducing computational time, compared to the high-fidelity physics-based numerical simulation (see Table 1). One example is a surrogate model based on surge response functions (SRFs) that are physics-based algebraic expressions of storm surge at a specific location of interest (Irish et al., 2009; Song et al., 2012; Taylor et al., 2015). Moreover, there are several other surrogate modeling approaches such as a multilayer perceptron (MLP, a class of feedforward artificial neural network), kriging (also referred to as Gaussian process regression), and support vector regression (SVR). The performance of these approaches are discussed in several studies (e.g., Hashemi et al., 2016; Al Kajbaf and Bensi, 2020). However, these surrogate models are designed only to predict one or two locations at a time and did not consider overland locations, which are too inefficient to apply across an extensive coastal region that spans a large number of geographic locations. For these reasons, several other surrogate modeling approaches have been developed, such as a moving least squares response surface (Taflanidis et al., 2013a,b), kriging with principal component analysis (PCA) (Jia and Taflanidis, 2013; Jia et al., 2016; Bass and Bedient, 2018), kriging with PCA and k-means clustering (Kyprioti et al., 2021), and an MLP with PCA (Bass and Bedient, 2018), to predict storm surges at a large number of geographic locations at the same time. Moreover, to consider overland locations, several methods were developed to estimate storm surge at the overland location from the water level at the offshore location using a one-dimensional Boussinesq model (Taflanidis et al., 2013b), from the water level at the nearest location (Taflanidis et al., 2013a), or from the water level at the nearby locations using k-nearest neighbor interpolation (Kyprioti et al., 2021).

The aforementioned surrogate modeling studies have three limitations. First, they did not consider a tropical cyclone (TC) track time-series but only considered a TC condition at one moment (usually at landfall) as model input (see Table 1). This type of point-based input, which does not take advantage of the temporal variation in TC conditions, limits the surrogate model's ability to distinguish storms with the same conditions at landfall but with different conditions before and/or after landfall. The implications of this simplified approach to characterize the storm track are potentially significant, given that historical records indicate that peak storm surge is not always most correlated with TC parameters at landfall (Needham and Keim, 2014). Secondly, several surrogate modeling studies did not consider bypassing storms that pass by the study area but do not make landfall within or near the study area (see Table 1). For example, Hashemi et al. (2016) and Kyprioti et al. (2021) used the North Atlantic Comprehensive Coastal Study (NACCS) synthetic tropical cyclone surge database of the US Army Corps of Engineers (Cialone et al., 2015), which contains storm surge simulations for 1050 synthetic tropical cyclones (TCs) in the US North Atlantic region (Virginia
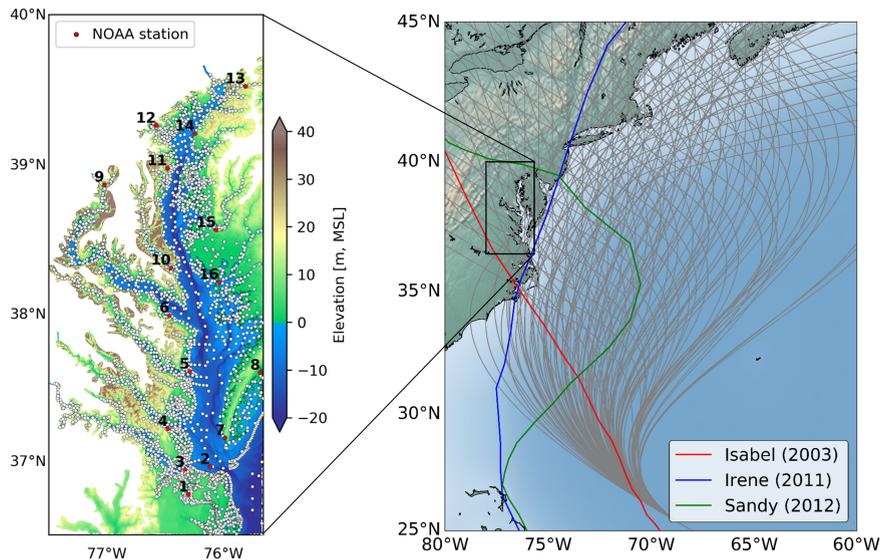
Figure 1: Location map of the Chesapeake Bay, Virginia and Maryland, USA. In the left topographic map (data from Cialone et al. (2015) and references therein), the white circles are the NACCS save points (Cialone et al., 2015), and the red circles represent the NOAA stations (Tides and Currents, 2021) (see Table 2). In the right map, the gray lines are the NACCS synthetic tropical cyclone (TC) tracks. The red, blue, and green lines represent the historical storm tracks of Hurricane Isabel in 2003, Hurricane Irene in 2011, and Hurricane Sandy in 2012, respectively.

to Maine). The NACCS defined the bypassing storm as a storm where track headings is larger than zero degree clockwise from North. Although bypassing storms make up approximately 40% of the NACCS database, Hashemi et al. (2016) and Kyprioti et al. (2021) excluded bypassing storms in surrogate modeling. The lack of capacity to respond to bypassing-track storms is a major problem for surrogate models because TCs, such as Hurricane Irene (2011), cannot be addressed. Thirdly, most of the aforementioned surrogate models' performance was only evaluated using cross-validation based on synthetic TCs (see Table 1). The cross-validation approach is often used to evaluate a surrogate model's performance over the entire database because this approach does not require new observations (Jia et al., 2016). However, especially when using a high-fidelity storm surge database where the synthetic storms are selected to optimize hazard characterization rather than surrogate model development (e.g., IPET, 2006; Kennedy et al., 2012; Cialone et al., 2015), the cross-validation approach may exaggerate the performance of the surrogate model. This is because performance is evaluated only with storms similar to those already used in developing the surrogate model (e.g., a storm that has the same storm track and a small difference in other parameter values)(Al Kajbaf and Bensi, 2020).

This study introduces a new surrogate model that can rapidly predict peak storm surges across an extensive coastal region from TC track time-series: hereafter referred to as the C1PKNet model, which stands for one-dimensional convolutional neural network model combined with principal component analysis and k-means clustering. Since the C1PKNet model uses the TC track time-series as model input, the C1PKNet model can predict a storm surge for a more realistic TC track than the existing surrogate models that only considered a TC condition at one moment as model input. This study's main contributions are as follows: (1) Our work is the first attempt to incorporate the TC track time-series in storm surge surrogate modeling. (2) We included bypassing storms, which are often neglected in surrogate modeling. (3) We not only evaluated the performance of the C1PKNet model based on a cross-validation approach but also validated the C1PKNet model with historical peak storm surge observations.

The remainder of the paper is organized as follows. The details regarding study area, data, and model development are provided in Section 2. The performance of the C1PKNet model based on synthetic tropical cyclones as well as historical hurricanes is provided in Section 3. The results of error contribution analysis are provided in Section 4. Future work is discussed in Section 5. The summary and the recommendations are provided in Section 6.

Table 2
NOAA Station ID and locations (data from Tides and Currents (2021))

No.	NAME	Station ID	LAT(°)	LON(°)
1	Money Point, VA	8639348	36.7783	-76.3017
2	Chesapeake Bay Bridge Tunnel, VA	8638863	36.9667	-76.1133
3	Sewells Point, VA	8638610	36.9467	-76.3300
4	Yorktown USCG Training Center, VA	8637689	37.2267	-76.4783
5	Windmill Point, VA	8636580	37.6167	-76.2900
6	Lewisetta, VA	8635750	37.9950	-76.4650
7	Kiptopeke, VA	8632200	37.1650	-75.9883
8	Wachapreague, VA	8631044	37.6083	-75.6850
9	Washington, DC	8594900	38.8733	-77.0217
10	Solomons Island, MD	8577330	38.3167	-76.4500
11	Annapolis, MD	8575512	38.9833	-76.4817
12	Baltimore, MD	8574680	39.2667	-76.5800
13	Chesapeake City, MD	8573927	39.5267	-75.8100
14	Tolchester Beach, MD	8573364	39.2133	-76.2450
15	Cambridge, MD	8571892	38.5717	-76.0617
16	Bishops Head, MD	8571421	38.2200	-76.0383

2. Methods

2.1. Study area

The Chesapeake Bay, the largest estuary in the United States, is one of the most complex coastal plain estuaries in the world, consisting of numerous tributaries, embayments, marshes, islands, and channels (Shen et al., 2006) (Figure 1). This mid-Atlantic area is vulnerable to flooding due to storm surges, as there are many geographically low-lying areas in the coastal plains surrounding the Chesapeake Bay (Li et al., 2020). Hurricane Isabel (2003) is a prime example that affected the Chesapeake Bay, where storm surge reached 2.5 m in the upper Chesapeake Bay, causing about $3.37 billion in losses (NOAA, 2004; Sheng et al., 2010). Even though observed storm surges in the Chesapeake from Hurricane Irene (2011) and Hurricane Sandy (2012) were not as large as the storm surges from Hurricane Isabel (2003), forecasted storm surges for these two hurricanes were large enough to issue voluntary and mandatory evacuations for low-lying and oceanfront areas around the Chesapeake Bay. Given the unique characteristics of the Chesapeake Bay, it provides a useful testbed to challenge the capabilities of machine-learning derived surrogate models.

2.2. Data

The C1PKNet model was developed based on the North Atlantic Comprehensive Coastal Study (NACCS) synthetic tropical cyclone surge database of the US Army Corps of Engineers (Cialone et al., 2015). The NACCS database contains storm surge simulations for 1050 synthetic tropical cyclones (TCs) in the US North Atlantic region (Virginia to Maine), computed using a coupled high-fidelity physics-based numerical modeling framework composed of ADCIRC (Luettich et al., 1992, ADvanced CIRCulation) and STWAVE (Massey et al., 2011, STeady-state spectral WAVE model) (hereafter referred to as "ADCIRC simulations"). A planetary boundary layer (PBL) numerical model is used to generate wind and pressure fields (Thompson and Cardone, 1996). The model grids consist of 3.1 million nodes and 6.2 million elements, where the smallest element's nodal spacing is approximately 10 m. Time-series of model results are saved for a subset of points (18,977 save-point locations). The 1050 synthetic TCs from the NACCS were selected to optimize hazard characterization via the joint probability method with optimized sampling (JPM-OS) (Resio et al., 2007; Toro et al., 2010). These synthetic TCs include both landfalling storms and bypassing storms. Each synthetic storm track is associated with a geographic reference location (Cialone et al., 2015). For landfalling storms, the NACCS defined the reference location as a point located 250 km before landfall. For bypassing storms, the NACCS defined the reference location as a point when the storm crossed a specific latitude. In this study, we used the 1031 synthetic TCs for which NACCS high-fidelity surge-only (no astronomical tide) simulations are available; due to ADCIRC modeling instabilities, simulated surges were not reported for 19 storms (Cialone, personal communication) (see gray lines in Figure 1). More details on the high-fidelity modeling and storm selection can be found in Cialone et al. (2015), Nadal-Caraballo et al. (2015) and Cialone et al. (2017). Note that we used the storm and save-point location indexing of
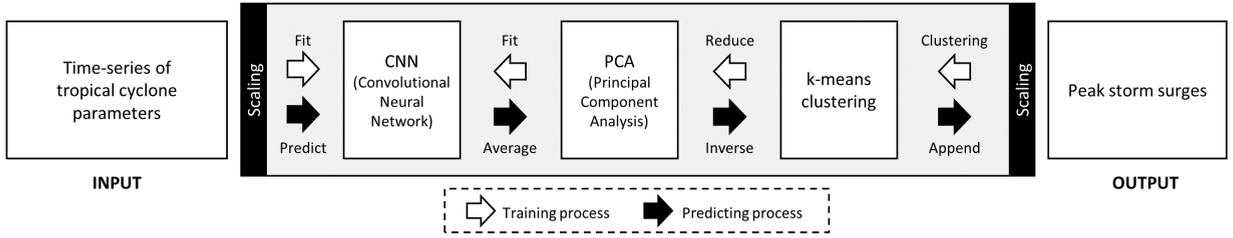
Figure 2: Computational flow of C1PKNet model.

The inputs of the C1PKNet model are the time-series of six TC parameters (latitude LAT, longitude LON, heading direction θ, central pressure C_p, radius of maximum winds R_{max}, and translation speed V_f)(Figure 2). The NACCS provides two other TC parameters (Holland B and a scale pressure radius) that are correlated with other TC parameters. Holland B is a function of the radius of maximum winds and latitude of storm location, and the scale pressure radius is a function of the radius of maximum winds. Thus, we did not use the Holland B and the scale pressure radius in this study in order to make the training process faster. To define the portion of the synthetic storm track time-series used for model development, we extracted the portion of the track beginning 30 hours before passing the reference point and up to 9 hours after passing the reference point, with track parameters specified at a 1-hour interval. This 39-hour time range is the maximum possible range that can be consistently extracted from all 1031 synthetic storm tracks.

The outputs of the C1PKNet model are the peak storm surges (Figure 2). To train the C1PKNet model, we used the NACCS-reported peak storm surges, referenced to Mean Sea Level (MSL), at 3,111 save-point locations within the study area (see white circles in Figure 1); NACCS save points that did not flood in any ADCIRC simulation are excluded. Out of these 3,111 locations, 21% of locations are overland locations (where the land elevation is higher than the MSL), and 79% of locations are subaqueous locations (where the land elevation is lower than the MSL). Among the entire dataset (e.g., in our case, all 3.2M values), about 20% of the NACCS-reported peak storm surges were missing for two reasons. First, about 3% of data were missing and marked by the U.S. Army Corps with the flag value because the data were not considered reliable due to the ADCIRC modeling instabilities. Secondly, about 17% of the data were missing because some of the overland locations were not inundated by storms. Training a surrogate model using a dataset with many missing values can have a significant impact on the performance of the surrogate model. To handle missing data, similar to Taflanidis et al. (2013a)'s approach, we filled the missing data at the overland locations with the peak storm surge of the nearest inundated overland location. If the surrounding ten overland locations were not inundated, we used the overland elevation to fill the missing value. Similarly, we filled the missing data at subaqueous locations with the nearest subaqueous location's peak storm surge. This complete dataset without missing data for 1031 synthetic TCs and 3,111 save-point locations was used for C1PKNet model training.

2.3. Model

Figure 2 shows the computational flow of the C1PKNet model. The inputs are the time-series of six TC parameters (LAT, LON, θ, C_p, R_{max}, V_f). The outputs are the peak storm surges at the 3,111 save-point locations (see white circles in Figure 1). The C1PKNet model is trained based on the synthetic TC track time series peak storm surges at all locations by following the direction of the white arrows in Figure 2. Once trained, the C1PKNet model can then predict, within a few seconds, the peak storm surges at all locations from any TC track time series by following the direction of the black arrows in Figure 2.

The C1PKNet model consists of three components: (1) k-means clustering, (2) principal component analysis (PCA), and (3) one-dimensional convolutional neural network (CNN). First, the k-means clustering groups the high-dimensional output (peak storm surges) among geospatial points with similar peak storm surge responses in order to increase model accuracy. Secondly, the PCA maps the peak storm surge responses of each cluster into a latent space and then reduces the dimension before CNN training, again in order to increase model efficiency. And lastly, the variation of TC conditions over time is incorporated into the C1PKNet model by applying a convolutional neural network (CNN), a class of deep neural networks widely used in several classification and regression applications. The main advantage of CNN is that it can capture spatial and/or temporal dependencies by preserving data structures in a process
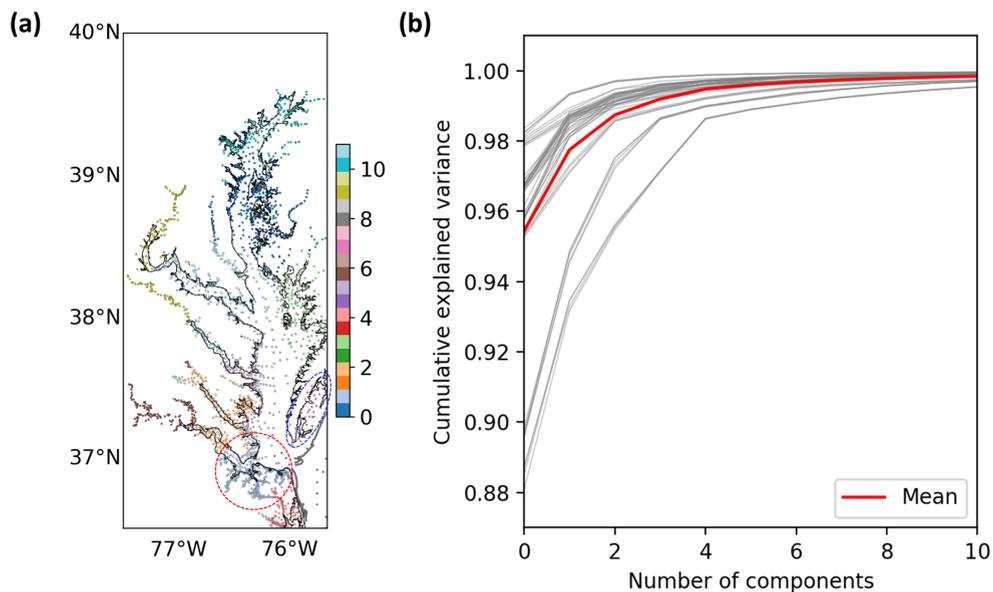
Figure 3: (a) Spatial distribution of clusters for the first outer-fold. Each color represents a different cluster. The red dashed circle is the Hampton-Norfolk area, and the blue dashed ellipse is the Hog Island bay-Mockhorn bay area. (b) Variation of cumulative explained variance versus the number of principal components. The gray lines represent the results according to each outer-fold and cluster, while the red line is the mean value.

called feature learning. Since the input TC conditions are time-series data, a one-dimensional CNN where the kernel (or filter) slides along the time dimension is used to capture the features of TC conditions over time. The C1PKNet model is written in Python where the k-means clustering and PCA were implemented using scikit-learn (Pedregosa et al., 2011), and the CNNs were implemented using Keras (Chollet et al., 2015) with a TensorFlow backend.

To prevent biased model evaluation, we trained and validated the C1PKNet model using nested k-fold cross-validation with shuffling (Figure S1). We set outer 10-fold cross-validation by separating 10% of TCs as a test set. At every iteration of the outer fold, inner 5-fold cross-validation was conducted where 20% of TCs were separated as a validation set. The entire process (e.g., k-means clustering, PCA, and CNN) was repeated for each outer-fold.

The training process is as follows (see white arrows in Figure 2). First, the inputs and outputs are standardized to reduce training time and improve model accuracy. Each of the inputs (six TC parameter time-series) was standardized by subtracting the mean and then scaling to unit variance based only on training input data. The outputs (the NACCS-reported peak storm surges at 3,111 save-point locations) were standardized by subtracting the mean and then scaling to unit variance based only on training output data. Note that these mean and variance values were saved in order to be used in the prediction process. Secondly, the standardized outputs are grouped into 12 clusters using k-means clustering. The number of clusters in k-means clustering was set by considering both the Chesapeake Bay's geographic features and the need for computational efficiency. One of the criteria was to separate the hydrodynamically distinct Hampton-Norfolk area (red dashed circle in Figure 3a) from the Hog Island bay-Mockhorn bay area (blue dashed circle in Figure 3a), and we found that at least 12 clusters are required to separate these two areas. The clusters for all 10 outer-folds can be found in Figure S2. Thirdly, for each cluster of each outer-fold, the PCA is conducted and retains five leading principal components. More details about the PCA can be found in Jolliffe and Cadima (2016). The number of principal components in PCA was set by considering the cumulative explained variance (Figure 3b). Our criterion was to keep the leading principal components, explaining, on average, more than 99% of the variance in peak storm surges in the clustered locations, and we found that at least five leading principal components are required to satisfy this criterion. Note that the eigenvector matrix of PCA was saved to reverse the principal components back to standardized peak storm surges in the predicting process (black arrows in Figure 2). And lastly, the CNNs were trained based on standardized inputs and five principal components. Figure 4 shows the detailed CNN architecture of the C1PKNet model. We use a batch normalization layer (Ioffe and Szegedy, 2015) after a one-dimensional convolutional layer in order to improve training efficiency and stability. We also use a dropout layer (Srivastava et al., 2014) after a
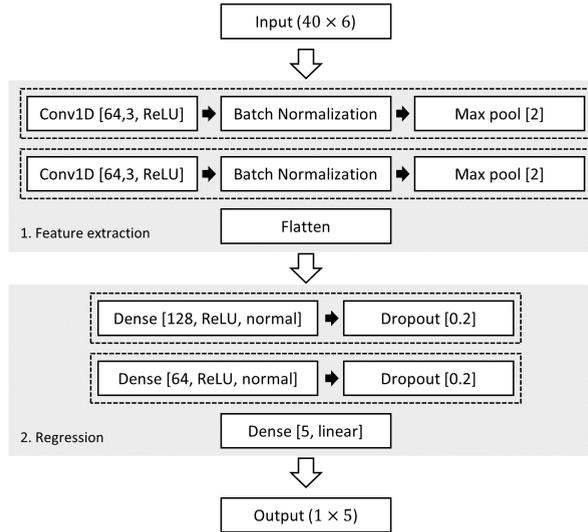
Figure 4: Convolutional neural network (CNN) architecture. Conv1D [64,3, ReLU] represents the 1-dimensional convolutional layer where the dimensionality of the output space (filter) is 64, the kernel size is 3, and the activation function is rectified linear activation function (ReLU). Max pool [2] represents the maximum pooling operation for 1-dimensional data, where the maximum pooling window size is 2. Dense [128, ReLU, normal] represents the densely-connected neural network layer, where the dimensionality of the output space is 128, the activation function is rectified linear activation function (ReLU), and the normal distribution is used as an initializer for the kernel weights matrix. Dropout [0.2] represents the dropout layer where the drop rate is 0.2. Dense [5, linear] represents the densely-connected neural network layer where the dimensionality of the output space is 5, and the activation function is a linear activation function.

dense layer to prevent over-fitting during training. In this study, we propose hyperparameters–namely, the number of layers, filter size, kernel size, pool size, number of neurons–as shown in Figure 4; hyperparameters were determined by considering both accuracy and efficiency. We used the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.01 and mean squared error (MSE) as a loss function. The maximum epoch was set to 10,000, and the CNN model training was set to stop early when the MSE of the validation set is not reduced after 100 epochs. The best CNN, namely the one that yields the minimum MSE, for the validation set was saved as the final trained CNN. Note that we trained 600 CNNs (10 outer folds * 5 inner folds *12 clusters) in this study.

Once the C1PKNet model is trained, the peak storm surges at 3,111 save-point locations can be predicted from any TC track time series as follows (see black arrows in Figure 2). First, the TC track time series for the case in which we want to predict storm surge was standardized using the mean and the variance of training input data. Second, five principal components were predicted by inputting the standardized TC track time series to the CNNs corresponding to each cluster. The average of five inner-fold estimates was considered as the final output of the CNNs corresponding to each cluster. Third, the predicted principal components were converted to the standardized peak storm surges using the eigenvector matrix of PCA derived in the training process. Fourth, the standardized predicted peak storm surges at the 3,111 save-point locations were scaled back to dimensional values using the mean and the variance of training output data. And lastly, the wet/dry conditions are checked by comparing the predicted peak storm surge and the land elevation. If the predicted peak storm surge is larger than the land elevation, the corresponding location is considered as a wet point. On the other hand, if the predicted peak storm surge is smaller than the land elevation, the corresponding location is considered as a dry point.

3. Performance test

3.1. Synthetic tropical cyclones

To quantify the performance of the C1PKNet model based on 1031 synthetic tropical cyclones (TCs), we used the 10-fold cross-validation method as mentioned in Section 2.3. For each outer fold, the 60 CNNs (5 inner folds × 12 clusters) were trained, and the C1PKNet model predicted peak storm surges for the test set that was not used in training.

This process was repeated 10 times to test the entire 1031 synthetic TCs. Here, we will first present the detailed results based on one of the synthetic TCs (Figures 5 and 6). Then, we will highlight the overall performance based on the locations where the wet/dry condition predicted by the C1PKNet model match with the ADCIRC simulation (Figure 7) and on the locations where the wet/dry condition predicted by the C1PKNet model does not match with the ADCIRC simulation results (Figure 8).

In this study, the error of the C1PKNet model ($\Delta\eta_n^i$) for storm i is represented as follows:

$$\Delta\eta_n^i = \eta_{C1PKNet,n}^i - \eta_{ADCIRC,n}^i \tag{1}$$

where i is the storm number, n is the point number, $\eta_{C1PKNet,n}^i$ is the peak storm surge predicted by the C1PKNet model, and $\eta_{NACCS,n}^i$ is the ADCIRC-simulated peak storm surge. The performance of the C1PKNet for each synthetic TC is represented by a root mean square error ($RMSE$) and mean bias error (MBE):

$$RMSE^i = \sqrt{\frac{1}{N_p} \sum_{n=1}^{N_p} \left[\Delta\eta_n^i\right]^2} \tag{2}$$

$$MBE^i = \frac{1}{N_p} \sum_{n=1}^{N_p} \Delta\eta_n^i \tag{3}$$

where N_p is the number of save-point locations considered in the error statistics.

Figure 5 shows an example of synthetic TC storm scenario inputs (Scenario 109), for a TC that was not used in C1PKNet model training. This TC's parameters are similar to Hurricane Isabel in 2003. The TC passed the reference point (yellow circle in Figure 5) near Cape Lookout (North Carolina, USA), with a central pressure (C_p) of 945 hPa, translation speed (V_f) of 32 km/h, heading direction (θ) of $-20°$, and radius of maximum winds (R_{max}) of 113 km.

Figure 6 shows the comparison between the peak storm surges predicted by the C1PKNet model and the ADCIRC-simulated peak storm surges. Overall, the peak storm surges predicted by the C1PKNet model agree well with the ADCIRC-simulated peak storm surges ($RMSE^{109} = 0.17\ m$). The C1PKNet model slightly underestimated the peak storm surges, especially at the northern locations ($MBE^{109} = -0.09\ m$). The wet/dry conditions predicted by the C1PKNet model match the ADCIRC simulation result at 98% of the locations; just 14 locations (2% of the overland locations) were inconsistently classified as wet/dry (see triangles in Figure 6).

Fig. 7a shows the ADCIRC-simulated peak storm surge (x-axis) versus the error in peak storm surge estimated using the C1PKNet model (y-axis) for all 1031 synthetic storm scenarios. Figure 7d and Figure 7g show the errors based only on the subaqueous locations and the overland locations, respectively. Note that we only plot the data that are predicted as wet by both the C1PKNet model simulation and the ADCIRC simulation. Overall, the peak storm surge errors ($\Delta\eta$) across surge magnitude are aligned well with the zero error line, with a slight trend towards under-prediction by the C1PKNet model for larger ADCIRC-simulated peak surge values. The largest error ($-6.3\ m$) occurred at the subaqueous location (point 13256) by synthetic scenario 50, where the peak storm surge (η_{ADCIRC}) was 7.02 m. We believe that this largest error reflects numerical instabilities in the ADCIRC simulation, rather than error in the C1PKNet model, because the peak storm surge at point 13256 by synthetic scenario 50 was distinctively larger compared to that of other similar TCs (i.e., synthetic scenarios $51-56$ where the peak storm surge range is 0.54 $m - 1.01\ m$) that have the same storm track and a small difference in other parameter values (Figure S3).

Since most of the peak storm surge values are concentrated in the range of 0 m and 2 m as shown in Figure 7a,d and g, an aggregate $RMSE$ and MBE based on the entire data can be misleading in terms of the actual performance of a surrogate storm surge model for events of significance to communities (Al Kajbaf and Bensi, 2020). Therefore, we calculated the $RMSE$ of peak storm surges for individual storms as the lower bound truncation threshold increases from 0 m to 1.5 m, with 0.5 m intervals (Figure 7b). Figure 7e and Figure 7f show the $RMSE$ based only on the subaqueous locations and the overland locations, respectively. As the lower bound truncation threshold increases, the median of $RMSE$ increases, and the range of $RMSE$ also increases. We also calculated the MBE of peak storm surges for individual storms as the lower bound truncation threshold increases from 0 m to 1.5 m, with 0.5 m
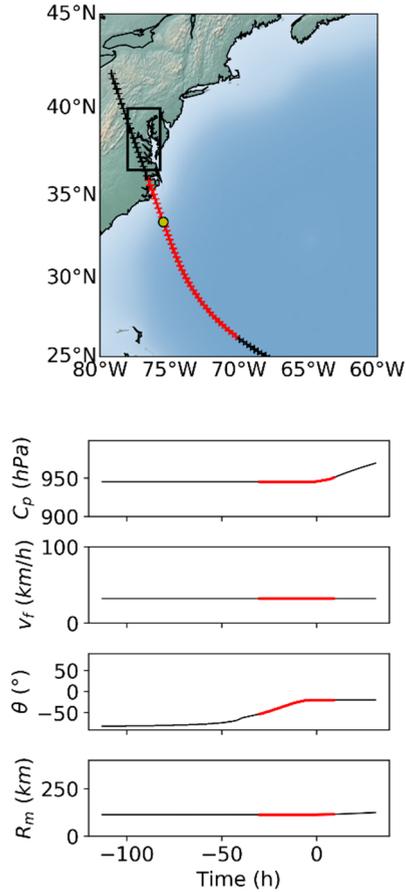
Figure 5: Example time series of synthetic tropical cyclone (Scenario 109) parameters where C_p is the central pressure, V_f is the translation speed, θ is the heading direction, and R_{max} is the radius of maximum winds. The yellow circle in the map represents the reference point that is 250 km away from the landfall point. The red lines are the 39 hours of TC parameters used as inputs in the C1PKNet model. The time of zero is set to the time when the TC passes the reference point.

intervals (Figure 7c). Figure 7f and Figure 7i show the MBE based only on the subaqueous locations and the overland locations, respectively. As the lower bound truncation threshold increases, the median of MBE decreases, and the range of MBE increases. Table 3 summarizes the $RMSE$ and MBE as well as how many scenarios are within the thresholds.

Figure 8 shows how consistent the C1PKNet model predictions of the wet/dry condition are with those of the ADCIRC simulations. First, for overland locations, we calculated the percentage of TC scenarios where the C1PKNet model's prediction of the wet/dry condition is the opposite of that predicted by ADCIRC simulation. Figure 8a shows the percentage of TC scenarios where each point is inundated according to the C1PKNet model prediction but not by the ADCIRC simulation (hereafter case 1). Figure 8b is shows the percentage of TC scenarios where each point point is inundated by the ADCIRC simulation but not by the C1PKNet model prediction (hereafter case 2). Overall, the C1PKNet model's prediction of the wet/dry condition differs from the ADCIRC prediction only for a few TC scenarios at each location. In case 1, considering all overland locations, the median of percentage of the mis-specified scenarios is 0.3%, and the maximum percentage of the mis-specified scenarios is 6.9%. In case 2, considering all overland locations, the median of the percentage of mis-specified scenarios is 0.4%, and the maximum percentage of mis-specified scenarios is 5.4%. In both cases, most of the overland locations where the wet/dry prediction is the opposite of that from ADCIRC were located in the low-lying area of coastal Virginia (below $37.5°N$). Secondly, for each synthetic TC scenario, we calculated the percentage of overland locations where the C1PKNet-predicted wet/dry
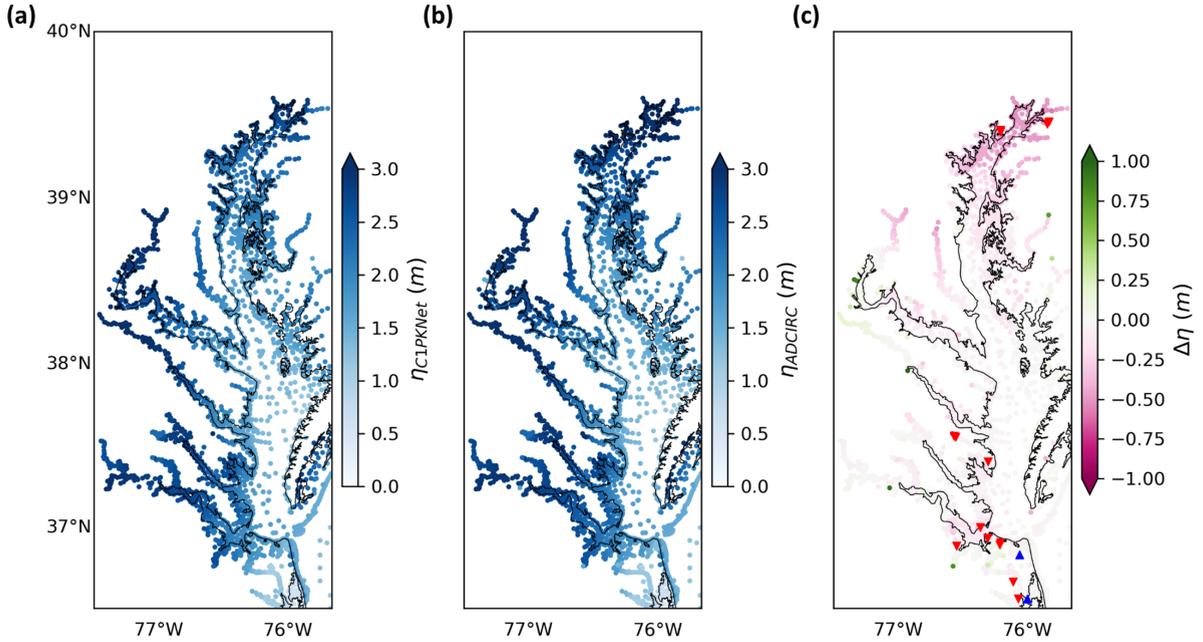
Figure 6: Synthetic tropical cyclone (TC) scenario result for TC shown in Figure 5. (a) Peak storm surges predicted by the C1PKNet model. (b) ADCIRC-simulated peak storm surges from the NACCS database. (c) The difference between the C1PKNet predictions and the NACCS predictions. The red triangles represent the locations where the C1PKNet model predicts that the location is not inundated and the ADCIRC model predicts that the location is inundated. The blue triangles represent the locations where the C1PKNet model predicts that the location is inundated and the ADCIRC model predicts that the location is not inundated.

Table 3
The number of tropical cyclone (TC) scenarios, root mean square error (RMSE), and mean bias error (MBE) depending on the lower bound truncation threshold.

	Threshold	TC scenario Count	RMSE (m) min	25%	50%	75%	max	MBE (m) min	25%	50%	75%	max
All locations	0	1031	0.01	0.02	0.03	0.06	0.54	-0.45	0.00	0.01	0.02	0.25
	0.5	726	0.00	0.03	0.06	0.10	0.55	-0.47	-0.04	-0.01	0.01	0.28
	1	424	0.02	0.07	0.11	0.17	0.78	-0.74	-0.11	-0.04	0.00	0.40
	1.5	308	0.03	0.09	0.14	0.24	1.32	-1.09	-0.15	-0.07	-0.01	0.41
Subaqueous locations	0	1031	0.01	0.02	0.03	0.06	0.55	-0.45	0.00	0.01	0.02	0.24
	0.5	726	0.00	0.03	0.06	0.10	0.55	-0.47	-0.04	-0.01	0.01	0.27
	1	421	0.00	0.06	0.10	0.17	0.78	-0.74	-0.11	-0.04	0.00	0.39
	1.5	303	0.01	0.08	0.13	0.23	1.44	-1.09	-0.15	-0.06	-0.01	0.41
Overland locations	0	1031	0.01	0.03	0.04	0.09	0.54	-0.45	0.00	0.01	0.02	0.34
	0.5	695	0.00	0.03	0.06	0.13	0.54	-0.46	-0.04	-0.01	0.01	0.35
	1	403	0.01	0.08	0.12	0.18	0.87	-0.84	-0.09	-0.04	0.00	0.43
	1.5	299	0.03	0.11	0.15	0.23	1.07	-1.07	-0.14	-0.07	-0.01	0.42

condition was opposite that predicted by ADCIRC (Figures 8c and d). For more than 96% of TC scenarios, per storm, there were less than 3% of overland locations mis-specified as wet or dry relative to the ADCIRC simulation. The median of case 1 is 1.1% of overland locations, while that of case 2 is 0.3% of overland locations. Lastly, we plotted the water depth (flood elevation less topographic elevation) frequency when the C1PKNet model's wet/dry prediction differed from the ADCIRC prediction (Figures 8e and f). The results show that the water depth was less than 0.5 *m* in more than 98% of the instances where the C1PKNet model's wet/dry predicting differs from that of the ADCIRC condition. For case 1, the water depth was less than 1 *m* for all data where C1PKNet predicted inundation and ADCIRC did not. For case 2, there were 16 occurrences where the water depth was larger than 1 *m*. In summary, the
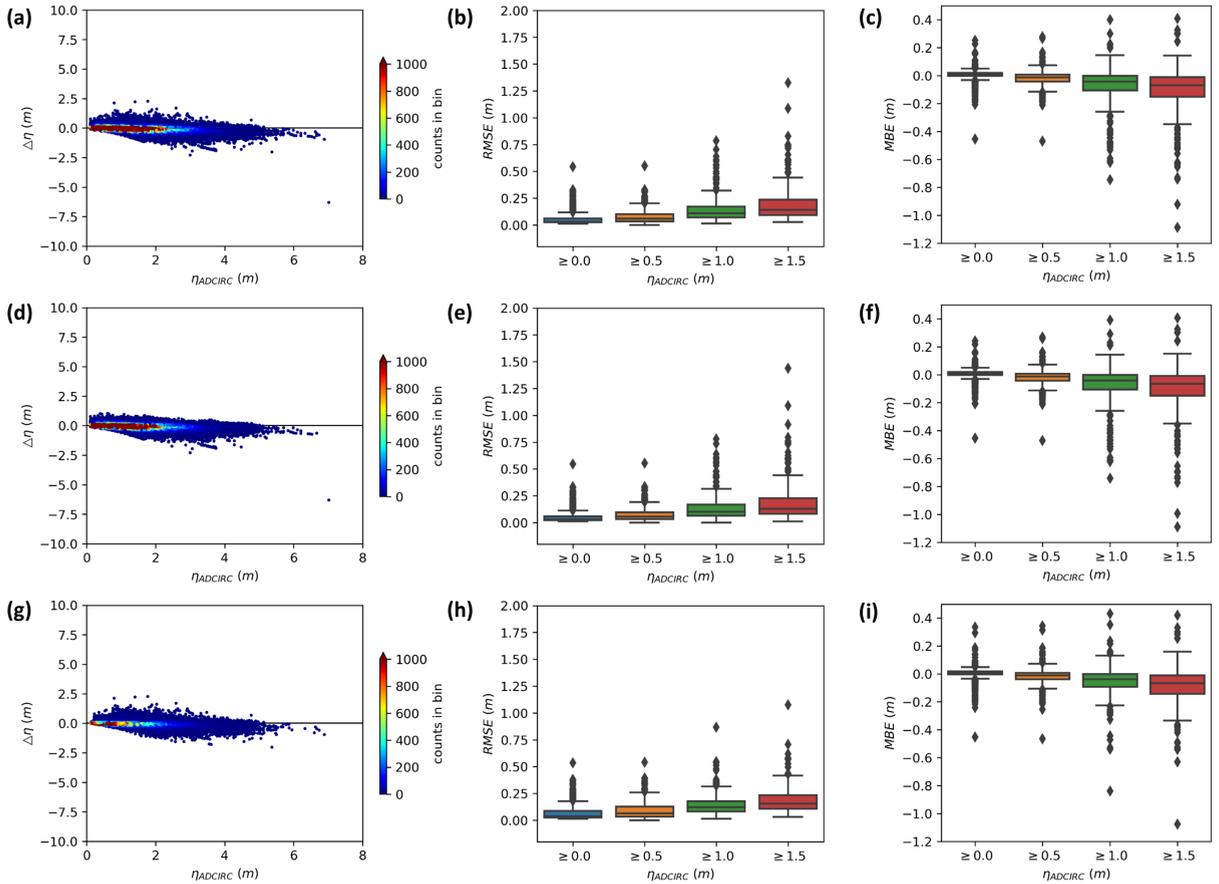
Figure 7: Performance of the C1PKNet model based on 1,031 synthetic storm scenarios. (Left panes) The density scatter plot that shows the error distribution depending on the ADCIRC-simulated peak storm surge from the NACCS database. The bin size of each axis is 5 cm. The black horizontal line represents the zero error line. (Middle panes) Boxplot of root mean square error (RMSE) of peak storm surges for individual storms as the lower bound truncation threshold increases from 0 m to 1.5 m. (Right panes) Boxplot of mean bias error (MBE) of peak storm surges for individual storms as the lower bound truncation threshold increases from 0 m to 1.5 m. The box symbol shows the interquartile range (box boundary) and median (horizontal line). The lower(upper) whisker is defined as 1.5 times the interquartile range below (above) the first (third) quartile. The data beyond the whiskers are classified as outliers (diamond). (a, b, and c) All locations. (d, e, and f) Subaqueous locations. (g, h, and i) Overland locations.

C1PKNet model predicted the wet/dry condition well, considering the number of storms and the number of locations. Moreover, even when the C1PKNet model predictions of wet/dry condition were not consistent with those of the ADCIRC simulations, the water depth was smaller than 0.5 m in most cases.

3.2. Historical hurricanes

Next, we tested the C1PKNet model based on three historical hurricanes: 2003 Isabel, 2011 Irene, and 2012 Sandy. We chose these three hurricanes because each of them shows different hurricane track features. Hurricane Isabel (2003) made landfall south of the study area (Outer Banks, North Carolina), Hurricane Irene (2011) bypasses the southeastern part of the study area, and Hurricane Sandy (2012) made landfall north of the study area (Brigantine, New Jersey). No recent hurricane has made direct landfall in the study area. Note that the existing surrogate models have generally not been developed to predict bypassing storms such as Hurricane Irene.

The hurricane tracks and the time-series of hurricane parameters are shown in Figure 9. Six-hourly hurricane-parameter data were taken from the extended best track dataset (Demuth et al., 2006), which supplements the second-generation hurricane database (HURDAT2) of the National Hurricane Center (NHC) (Landsea and Franklin, 2013).

Using linear interpolation, hourly hurricane parameter values over 39 hours were used as the inputs to the C1PKNet model (red lines in Figure 9). Here, the reference point at hour zero was set to the interpolated track point closest to the reference point of the most similar synthetic NACCS track.

The C1PKNet model predictions of three historical hurricanes are shown in Figure 10. In this study, we used a cross-validation ensemble (the average of ten trained C1PKNet model predictions) as a final prediction to reduce the variance of predictions and reduce generalization error. The maximum value of the peak storm surges was 3.41 m (Hurricane Isabel), 1.88 m (Hurricane Irene), and 0.85 m (Hurricane Sandy), respectively. Among the 16 NOAA stations in the study area (Table 2), storm surges were observed at 12 stations (Hurricane Isabel), 10 stations (Hurricane Irene), and 15 stations (Hurricane Sandy), respectively (red circles in Figure 10).

Figure 11 shows the comparison between the C1PKNet model results and the NOAA stations' observations. The peak storm tides ($\bar{\eta}$) were estimated as follows:

$$\bar{\eta} = \eta_{C1PKNet} + \eta_{Tide} - \eta_{NLR} \tag{4}$$

where $\eta_{C1PKNet}$ is a peak storm surge predicted by the C1PKNet model, η_{Tide} is a predicted tide elevation, and η_{NLR} is a location-dependent astronomical tide bias calculated by the NACCS to account for errors associated with superposition of tides and surge water levels. Note that we set $\eta_{C1PKNet}$ to the peak storm surge at the save point located closest to the NOAA station. The predicted tide data (η_{Tide}) and observed storm tide data were extracted from the NOAA reports (Hovis et al., 2004; Fanelli and Fanelli, 2011; Fanelli et al., 2013). Observational data were converted from mean lower low water (MLLW) to mean sea level (MSL) to be consistent with the C1PKNet output (NOAA, 2020). The NACCS-reported astronomical tide bias (η_{NLR}) was determined using ADCIRC simulations (Nadal-Caraballo et al., 2020). Overall, the C1PKNet model predictions match well with the observations. The $RMSE$ taken across available gauges for each storm is 0.29 m (Hurricane Isabel), 0.25 m (Hurricane Irene), and 0.44 m (Hurricane Sandy), respectively. For Hurricane Isabel (2003), the largest deviation, underestimated by 0.65 m, occurred at the Chesapeake City station. For Hurricane Irene (2011), the C1PKNet model slightly underestimated the large storm tides (storm tides > 1.8 m), by 0.16 m on average, while overestimating the small storm tides (storm tides < 1.8 m), by 0.25 m on average. For Hurricane Sandy (2012), the C1PKNet model consistently underestimated the storm tides ($MBE = -0.40$ m), underestimating by as little as 0.08 m to as much as 0.68 m. We believe that the relatively large underestimation of storm tide shown in the Hurricane Sandy (2012) case is due to sudden changes in hurricane parameters near the time of landfall; such dramatic changes were not captured in the synthetic scenarios in the NACCS storm set and used for the C1PKNet model training.

4. Error contribution analysis

In section 3.1, the performance of the C1PKNet model is investigated based only on the final outputs of the C1PKNet model (peak storm surge predictions). However, errors can arise from all three C1PKNet components: the convolutional neural network (CNN), principal component analysis (PCA), and k-means clustering. To explore the impact of each component separately, we performed three additional analyses as follows.

First, to investigate the errors coming from the CNN, we compared the principal components derived from the scaled peak storm surges of ADCIRC simulations and the principal components predicted by the CNNs across 10 folds (Figure 12). Overall, the CNNs predict well the first principal component (PC1) that explains 95.4% of the variance of peak storm surges ($R^2 = 0.99$). However, the accuracy of the CNNs gradually decreased from the second principal component (2.3% of variance, $R^2 = 0.91$) to the fifth principal component (0.3% of variance, $R^2 = 0.76$).

Secondly, to investigate the errors coming from PCA, we compared the original peak storm surges of NACCS and the peak storm surges reconstructed using the first five principal components for each outer-fold. Figure 13 shows the first outer fold result. Comparative figures for all 10 outer-folds can be found in Figure S5. Overall, the reconstructed peak storm surges match well with the ADCIRC-simulated peak storm surges ($R^2 \approx 1.0$). The difference between the ADCIRC-simulated peak storm surges and the reconstructed peak storm surges was greater than 0.5 m for just 0.03% of data. The three largest differences occur at point 13256 (Figure S3) and point 15125 (Figure S4) when ADCIRC-simulated peak storm surge is greater than 6 m. We believe these outliers arise from numerical instabilities within the ADCIRC simulations, as mentioned in section 3.1 because these ADCIRC-simulated peak storm surges were distinctively larger than that of other similar TCs at the same two points.

And lastly, to investigate the errors coming from k-means clustering, we first examined how similar the k-means clustering results were during 10-fold cross-validation. To quantify the cluster similarity at a point-level, we calculated how often a pair of points were classified into the same cluster for all folds. And then, for each point, we counted the number of points where the frequency was zero (never grouped into the same cluster) and ten (always grouped into the same cluster). Based on the counted number (N), we defined the cluster similarity (CS), which indirectly represents how often the point was classified into the same cluster, as follows:

$$CS = \frac{N}{N_T - 1} \times 100 (\%) \tag{5}$$

where N_T is the total number of locations considered in the study (3111 locations). As shown in Figure 14a, most of the points were classified into the same cluster during 10-fold cross-validation (median $CS = 98.1\%$). The relatively small CS values ($CS < 80\%$) were shown in 141 locations (about 4.5%) that are mainly distributed in the middle of the bay ($37°N - 38.5°N$). To compare the accuracy of the C1PKNet model depending on the CS values, we plotted the error distribution after dividing the locations into two groups: (1) locations where CS is larger than 98.1% and (2) locations where CS is less than or equal to 98.1% (Figure 14b and c). The results show that there is no notable difference between the two groups. The small CS group ($CS \leq 98.1\%$) had slightly more data showing errors larger than 0.5 m (1055 data) compared to the large CS group ($CS > 98.1\%$, 870 data). These results suggest that the frequency of how often the location is classified into the same cluster does not significantly affect the C1PKNet model accuracy.

5. Future work

There are still several factors that need to be addressed in the future if training data can be obtained. First, the US Northeast is affected by both extratropical cyclones (e.g., Nor'easters) and tropical cyclones (Bakhtyar et al., 2018), but the C1PKNet model in its current form is only applicable to tropical cyclones. Even though tropical cyclones have traditionally been more intense than extratropical cyclones, Nor'easters are known to affect the Chesapeake Bay more frequently, last longer, and impact larger areas than tropical cyclones (Gong et al., 2009; Garzon et al., 2018). Therefore, future studies should expand surrogate storm surge modeling applicability to extratropical cyclones.

Secondly, the C1PKNet model in its current form is only able to predict peak storm surges. Historical hurricanes such as Ike (2008) have shown that dangerous forerunner surge preceding a peak storm surge can occur (Kennedy et al., 2011; Liu and Irish, 2019). Although large forerunner surges are rare, when they do occur, they have the capacity to impede evacuation. Several studies have pointed out predicting morphological change during storms is critical to protect coastal communities from damaging sediment transport, deposition, and erosion (e.g., Hu et al., 2018; Smallegan and Irish, 2017). Thus, the C1PKnet model should be expanded to predict not only the storm surge time-series to capture the forerunner surge but also the potential for morphological change during storms.

Thirdly, several factors that can exacerbate flooding when combined with storm surge–namely astronomical tide, sea-level change, and heavy rainfall–were not considered in this study. Several studies have shown that the water level can be higher or lower than the sum of individual contributions by these processes, due to nonlinear interactions with the meteorologically driven surge (e.g., Thomas et al., 2019). Rising sea levels due to climate change and regional processes such as land subsidence, which can trigger increases in flood frequency by storm surge, are a great concern to coastal communities (Woodruff et al., 2013). The Hampton Roads area in the lower Chesapeake Bay is especially vulnerable as the region is experiencing a high rate of relative sea-level rise (Sallenger et al., 2012; Ezer and Atkinson, 2014; Buzzanga et al., 2020). If storm surges and heavy rains occur simultaneously (or in close succession), the possibility of flooding is much greater than from storm surge alone (Wahl et al., 2015; Bass and Bedient, 2018). Thus, future studies are necessary to take into account the interaction between storm surge and astronomical tide, sea-level change, and heavy rainfall in surrogate coastal flood modeling.

Lastly, there are several parameters of the C1PKNet model that need to be optimized. For example, as inputs of the C1PKNet model, we used the longest time-series of TC parameters (39 hours) that can be extracted from the NACCS database. Since the CNN training time is related to the input size, the time-series length should be optimized considering both model accuracy and efficiency. Also, various CNN architectures (Figure 4), the number of principal components, and the number of clusters should be tested to optimize the performance of the C1PKNet model. However, since this hyperparameter/parameter tuning will be computationally expensive with a grid search approach (considering all possible combinations), an efficient optimization algorithm will be needed. Moreover, Taflanidis et al. (2013a)'s

approach adopted herein to fill the missing peak storm surge values in this study is sensitive to the peak storm surge of the nearest point. Thus, future studies should investigate several interpolation approaches (e.g., k-nearest neighbor interpolation with an inverse distance weighting (Kyprioti et al., 2021)) that can improve the performance of the C1PKNet model.

6. Conclusions

This study presents a new surrogate model for tropical cyclone surge called the C1PKNet model that can rapidly predict peak storm surges across an extensive coastal region from time-series of tropical cyclone (TC) conditions. A convolutional neural network (CNN) was implemented in the surrogate modeling to incorporate the variation of tropical cyclone conditions over time. Principal component analysis and k-means clustering were integrated to facilitate the high efficiency and accuracy of CNN. The C1PKNet model was trained for the Chesapeake Bay using 1031 synthetic tropical cyclones from the NACCS database, including landfalling storms and bypassing storms. The performance of the C1PKNet model was evaluated based on not only 1031 synthetic tropical cyclones of the NACCS database but also on observations from three historical hurricanes (Hurricane Isabel in 2003, Hurricane Irene in 2011, and Hurricane Sandy in 2012). To provide complete performance information, we analyzed the error of individual point data and the aggregate error metric after grouping the data by peak storm surge values (Figure 7). We also analyzed how consistent the C1PKNet model predictions of the wet/dry condition are, with respect to those of the ADCIRC simulations (Figure 8). Moreover, we analyzed the errors coming from each component of the C1PKNet model (Figures 12-14). The results indicate that the C1PKNet model is computationally efficient and can predict peak storm surges for a realistic TC track that changes over time. We believe that the C1PKNet model has the potential to enhance coastal resilience by providing rapid storm surge predictions. From this study, the following recommendations can be drawn for surrogate modeling to predict peak storm surges:

1. The use of CNN allows for consideration of the the time-series of TC parameters rather than being limited to values taken at a single representative time-step.
2. Bypassing storms that can lead to impactful surge values should not be neglected in surrogate modeling.
3. The performance assessment should include case studies based on historical hurricanes because performance quantification relying solely on cross-validation can be misleading in terms of the actual performance of the surrogate model.

Acknowledgments

References

Al Kajbaf, A., Bensi, M., 2020. Application of surrogate models in estimation of storm surge:a comparative assessment. Applied Soft Computing 91, 106184. doi:https://doi.org/10.1016/j.asoc.2020.106184.

Bakhtyar, R., Orton, P., Marsooli, R., Miller, J., 2018. Rapid wave modeling of severe historical extratropical cyclones off the northeastern united states. Ocean Engineering 159, 315–332. doi:https://doi.org/10.1016/j.oceaneng.2018.04.037.

Bass, B., Bedient, P., 2018. Surrogate modeling of joint flood risk across coastal watersheds. Journal of Hydrology 558, 159–173. doi:https://doi.org/10.1016/j.jhydrol.2018.01.014.

Blake, E.S., Landsea, C., Gibney, E.J., 2011. The deadliest, costliest, and most intense United States tropical cyclones from 1851 to 2010 (and other frequently requested hurricane facts). Technical Report NOAA Technical Memorandum NWS NHC-6. National Weather Service, National Hurricane Center. Miami, Florida. URL: https://www.nhc.noaa.gov/pdf/nws-nhc-6.pdf.

Buzzanga, B., Bekaert, D.P.S., Hamlington, B.D., Sangha, S.S., 2020. Toward sustained monitoring of subsidence at the coast using insar and gps: An application in hampton roads, virginia. Geophysical Research Letters 47, e2020GL090013. doi:https://doi.org/10.1029/2020GL090013.

Chollet, F., et al., 2015. Keras. https://keras.io.

Cialone, M.A., Grzegorzewski, A.S., Mark, D.J., Bryant, M.A., Massey, T.C., 2017. Coastal-storm model development and water-level validation for the north atlantic coast comprehensive study. Journal of Waterway, Port, Coastal, and Ocean Engineering 143, 04017031. doi:10.1061/(ASCE)WW.1943-5460.0000408.

Cialone, M.A., Massey, T.C., Anderson, M.E., Grzegorzewski, A.S., Jensen, R.E., Cialone, A., Mark, D.J., Pevey, K.C., Gunkel, B.L., McAlpin, T.O., 2015. North Atlantic Coast Comprehensive Study (NACCS) coastal storm model simulations: Waves and water levels. Technical Report ERCD/CHL TR-15-14. U.S. Army Engineer Research and Development Center, Coastal and Hydraulics Laboratory. Vicksburg, MS. URL: https://hdl.handle.net/11681/7339.

Demuth, J.L., DeMaria, M., Knaff, J.A., 2006. Improvement of advanced microwave sounding unit tropical cyclone intensity and size estimation algorithms. Journal of applied meteorology and climatology 45, 1573–1581.

Ezer, T., 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, 362–382. doi:https://doi.org/10.1002/2014EF000252.

Fanelli, C., Fanelli, P., 2011. Hurricane irene. URL: https://repository.library.noaa.gov/view/noaa/17128. Technical Report.

Fanelli, C., Fanelli, P., Wolcott, D., 2013. Hurricane sandy. URL: https://repository.library.noaa.gov/view/noaa/17125. Technical Report.

Garzon, J.L., Ferreira, C.M., Padilla-Hernandez, R., 2018. Evaluation of weather forecast systems for storm surge modeling in the chesapeake bay. Ocean Dynamics 68, 91–107. doi:10.1007/s10236-017-1120-x.

Glahn, B., Taylor, A., Kurkowski, N., Shaffer, W.A., 2009. The role of the slosh model in national weather service storm surge forecasting. National Weather Digest 33, 3–14.

Gong, W., Shen, J., Cho, K.H., Wang, H.V., 2009. A numerical model study of barotropic subtidal water exchange between estuary and subestuaries (tributaries) in the chesapeake bay during northeaster events. Ocean Modelling 26, 170–189. doi:https://doi.org/10.1016/j.ocemod.2008.09.005.

Hashemi, M.R., Spaulding, M.L., Shaw, A., Farhadi, H., Lewis, M., 2016. An efficient artificial intelligence model for prediction of tropical storm surge. Natural Hazards 82, 471–491. doi:10.1007/s11069-016-2193-4.

Hovis, G., Popovich, W., Zervas, Chris Eugene, ., Hubbard, J.R., Shih, H.H., Stone, P.J., 2004. Effects of hurricane isabel on water levels : data report. URL: https://repository.library.noaa.gov/view/noaa/14665. Technical Report.

Hu, K., Chen, Q., Wang, H., Hartig, E.K., Orton, P.M., 2018. Numerical modeling of salt marsh morphological change induced by hurricane sandy. Coastal Engineering 132, 63–81. doi:https://doi.org/10.1016/j.coastaleng.2017.11.001.

Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: International conference on machine learning, PMLR. pp. 448–456.

IPET, 2006. Performance evaluation of the New Orleans and southeast Louisiana hurricane protection system: Final report of the interagency performance evaluation task force. Technical Report.

Irish, J.L., Resio, D.T., Cialone, M.A., 2009. A surge response function approach to coastal hazard assessment. part 2: Quantification of spatial attributes of response functions. Natural Hazards 51, 183–205. doi:10.1007/s11069-009-9381-4.

Jelesnianski, C.P., 1992. SLOSH: Sea, lake, and overland surges from hurricanes. volume 48. US Department of Commerce, National Oceanic and Atmospheric Administration . . . .

Jia, G., Taflanidis, A.A., 2013. Kriging metamodeling for approximation of high-dimensional wave and surge responses in real-time storm/hurricane risk assessment. Computer Methods in Applied Mechanics and Engineering 261-262, 24–38. doi:https://doi.org/10.1016/j.cma.2013.03.012.

Jia, G., Taflanidis, A.A., Nadal-Caraballo, N.C., Melby, J.A., Kennedy, A.B., Smith, J.M., 2016. Surrogate modeling for peak or time-dependent storm surge prediction over an extended coastal region using an existing database of synthetic storms. Natural Hazards 81, 909–938. doi:10.1007/s11069-015-2111-1.

Jolliffe, I.T., Cadima, J., 2016. Principal component analysis: a review and recent developments. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 374, 20150202. doi:10.1098/rsta.2015.0202.

Kennedy, A.B., Gravois, U., Zachry, B.C., Westerink, J.J., Hope, M.E., Dietrich, J.C., Powell, M.D., Cox, A.T., Luettich Jr., R.A., Dean, R.G., 2011. Origin of the hurricane ike forerunner surge. Geophysical Research Letters 38. doi:https://doi.org/10.1029/2011GL047090.

Kennedy, A.B., Westerink, J.J., Smith, J.M., Hope, M.E., Hartman, M., Taflanidis, A.A., Tanaka, S., Westerink, H., Cheung, K.F., Smith, T., Hamann, M., Minamide, M., Ota, A., Dawson, C., 2012. Tropical cyclone inundation potential on the hawaiian islands of oahu and kauai. Ocean Modelling 52-53, 54–68. doi:https://doi.org/10.1016/j.ocemod.2012.04.009.

Kerr, P.C., Martyr, R.C., Donahue, A.S., Hope, M.E., Westerink, J.J., Luettich Jr., R.A., Kennedy, A.B., Dietrich, J.C., Dawson, C., Westerink, H.J., 2013. U.s. ioos coastal and ocean modeling testbed: Evaluation of tide, wave, and hurricane surge response sensitivities to mesh resolution and friction in the gulf of mexico. Journal of Geophysical Research: Oceans 118, 4633–4661. doi:10.1002/jgrc.20305.

Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 .

Krien, Y., Dudon, B., Roger, J., Zahibo, N., 2015. Probabilistic hurricane-induced storm surge hazard assessment in guadeloupe, lesser antilles. Natural Hazards and Earth System Sciences 15, 1711–1720. doi:10.5194/nhess-15-1711-2015.

Kyprioti, A.P., Taflanidis, A.A., Nadal-Caraballo, N.C., Campbell, M., 2021. Storm hazard analysis over extended geospatial grids utilizing surrogate models. Coastal Engineering , 103855doi:https://doi.org/10.1016/j.coastaleng.2021.103855.

Landsea, C.W., Franklin, J.L., 2013. Atlantic hurricane database uncertainty and presentation of a new database format. Monthly Weather Review 141, 3576–3592.

Li, M., Zhang, F., Barnes, S., Wang, X., 2020. Assessing storm surge impacts on coastal inundation due to climate change: case studies of baltimore

and dorchester county in maryland. Natural Hazards 103, 2561–2588. doi:10.1007/s11069-020-04096-4.

Lin, N., Emanuel, K.A., Smith, J.A., Vanmarcke, E., 2010. Risk assessment of hurricane storm surge for new york city. Journal of Geophysical Research: Atmospheres 115. doi:10.1029/2009JD013630.

Liu, Y., Irish, J.L., 2019. Characterization and prediction of tropical cyclone forerunner surge. Coastal Engineering 147, 34–42. doi:https://doi.org/10.1016/j.coastaleng.2019.01.005.

Luettich, R.A., Westerink, J.J., Scheffner, N.W., et al., 1992. ADCIRC: an advanced three-dimensional circulation model for shelves, coasts, and estuaries. Report 1, Theory and methodology of ADCIRC-2DD1 and ADCIRC-3DL. Technical Report. URL: https://hdl.handle.net/11681/4618.

Massey, T.C., Anderson, M.E., Smith, J.M., Gomez, J.M., Jones, R., 2011. Stwave: Steady-state spectral wave model user's manual for stwave, version 6.0 .

Nadal-Caraballo, N.C., Campbell, M.O., Gonzalez, V.M., Torres, M.J., Melby, J.A., Taflanidis, A.A., 2020. Coastal hazards system: A probabilistic coastal hazard analysis framework. Journal of Coastal Research 95, 1211–1216. doi:10.2112/SI95-235.1.

Nadal-Caraballo, N.C., Melby, J.A., Gonzalez, V.M., Cox, A.T., 2015. Coastal storm hazards from Virginia to Maine. Technical Report ERDC/CHL TR-15-5. U.S. Army Engineer Research and Development Center, Coastal and Hydraulics Laboratory. Vicksburg, MS. URL: https://hdl.handle.net/11681/7715.

Needham, H.F., Keim, B.D., 2014. Correlating storm surge heights with tropical cyclone winds at and before landfall. Earth Interactions 18, 1 – 26. doi:10.1175/2013EI000527.1.

NOAA, 2004. Service Assessment: Hurricane Isabel September 18-19, 2003. Technical Report. National Weather Service, Silver Spring, Maryland. URL: https://repository.library.noaa.gov/view/noaa/6471.

NOAA, 2020. Tides and currents. URL: https://tidesandcurrents.noaa.gov/.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research 12, 2825–2830.

Rappaport, E.N., 2014. Fatalities in the united states from atlantic tropical cyclones: New data and interpretation. Bulletin of the American Meteorological Society 95, 341 – 346. doi:10.1175/BAMS-D-12-00074.1.

Resio, D.T., Asher, T.G., Irish, J.L., 2017. The effects of natural structure on estimated tropical cyclone surge extremes. Natural Hazards 88, 1609–1637. doi:10.1007/s11069-017-2935-y.

Resio, D.T., Boc, S.J., Borgman, L.E., Cardone, V.J., Cox, A.W., Dally, W.R., Dean, R.G., Divoky, D., Hirsh, E., Irish, J.L., Levinson, D., Niedoroda, A.W., Powell, M.D., Ratcliff, J.J., Stutts, V., Suhayda, J.N., Toro, G.R., Vickery, P.J., Westerink, J.J., 2007. White paper on estimating hurricane inundation probabilities. Technical Report. U.S. Army Engineer Research and Development Center, Coastal and Hydraulics Laboratory. Vicksburg, MS. URL: https://hdl.handle.net/11681/22643.

Resio, D.T., Irish, J.L., 2015. Tropical cyclone storm surge risk. Current Climate Change Reports 1, 74–84. doi:10.1007/s40641-015-0011-9.

Robertson, I.N., Riggs, H.R., Yim, S.C., Young, Y.L., 2007. Lessons from hurricane katrina storm surge on bridges and buildings. Journal of Waterway, Port, Coastal, and Ocean Engineering 133, 463–483. doi:10.1061/(ASCE)0733-950X(2007)133:6(463).

Sallenger, A.H., Doran, K.S., Howd, P.A., 2012. Hotspot of accelerated sea-level rise on the atlantic coast of north america. Nature Climate Change 2, 884–888. doi:10.1038/nclimate1597.

Shen, J., Wang, H., Sisson, M., Gong, W., 2006. Storm tide simulation in the chesapeake bay using an unstructured grid model. Estuarine, Coastal and Shelf Science 68, 1–16. doi:https://doi.org/10.1016/j.ecss.2005.12.018.

Sheng, Y.P., Alymov, V., Paramygin, V.A., 2010. Simulation of storm surge, wave, currents, and inundation in the Outer Banks and Chesapeake Bay during Hurricane Isabel in 2003: The importance of waves. Journal of Geophysical Research: Oceans 115. doi:https://doi.org/10.1029/2009JC005402.

Smallegan, S.M., Irish, J.L., 2017. Barrier island morphological change by bay-side storm surge. Journal of Waterway, Port, Coastal, and Ocean Engineering 143, 04017025. doi:10.1061/(ASCE)WW.1943-5460.0000413.

Smith, A.B., 2020. U.S. Billion-dollar Weather and Climate Disasters, 1980 - present (NCEI Accession 0209268). NOAA National Centers for Environmental Information. Dataset. URL: https://www.ncdc.noaa.gov/billions/, doi:10.25921/stkw-7w73.

Song, Y.K., Irish, J.L., Udoh, I.E., 2012. Regional attributes of hurricane surge response functions for hazard assessment. Natural Hazards 64, 1475–1490. doi:10.1007/s11069-012-0309-z.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. The journal of machine learning research 15, 1929–1958.

Taflanidis, A.A., Jia, G., Kennedy, A.B., Smith, J.M., 2013a. Implementation/optimization of moving least squares response surfaces for approximation of hurricane/storm surge and wave responses. Natural Hazards 66, 955–983. doi:10.1007/s11069-012-0520-y.

Taflanidis, A.A., Kennedy, A.B., Westerink, J.J., Smith, J., Cheung, K.F., Hope, M., Tanaka, S., 2013b. Rapid assessment of wave and surge risk during landfalling hurricanes: Probabilistic approach. Journal of Waterway, Port, Coastal, and Ocean Engineering 139, 171–182. doi:10.1061/(ASCE)WW.1943-5460.0000178.

Taylor, N.R., Irish, J.L., Udoh, I.E., Bilskie, M.V., Hagen, S.C., 2015. Development and uncertainty quantification of hurricane surge response functions for hazard assessment in coastal bays. Natural Hazards 77, 1103–1123. doi:10.1007/s11069-015-1646-5.

Thomas, A., Dietrich, J., Asher, T., Bell, M., Blanton, B., Copeland, J., Cox, A., Dawson, C., Fleming, J., Luettich, R., 2019. Influence of storm timing and forward speed on tides and storm surge during hurricane matthew. Ocean Modelling 137, 1–19. doi:https://doi.org/10.1016/j.ocemod.2019.03.004.

Thompson, E.F., Cardone, V.J., 1996. Practical modeling of hurricane surface wind fields. Journal of Waterway, Port, Coastal, and Ocean Engineering 122, 195–205. doi:10.1061/(ASCE)0733-950X(1996)122:4(195).

Tides, Currents, 2021. NOAA NOA/CO-OPS, Silver Spring, Maryland. URL: http://tidesandcurrents.noaa.gov/index.shtml/.

Toro, G.R., Resio, D.T., Divoky, D., Niedoroda, A.W., Reed, C., 2010. Efficient joint-probability methods for hurricane surge frequency analysis.

Ocean Engineering 37, 125–134.

Wahl, T., Jain, S., Bender, J., Meyers, S.D., Luther, M.E., 2015. Increasing risk of compound flooding from storm surge and rainfall for major us cities. Nature Climate Change 5, 1093–1097. doi:`10.1038/nclimate2736`.

Woodruff, J.D., Irish, J.L., Camargo, S.J., 2013. Coastal flooding by tropical cyclones and sea-level rise. Nature 504, 44–52. doi:`10.1038/nature12855`.
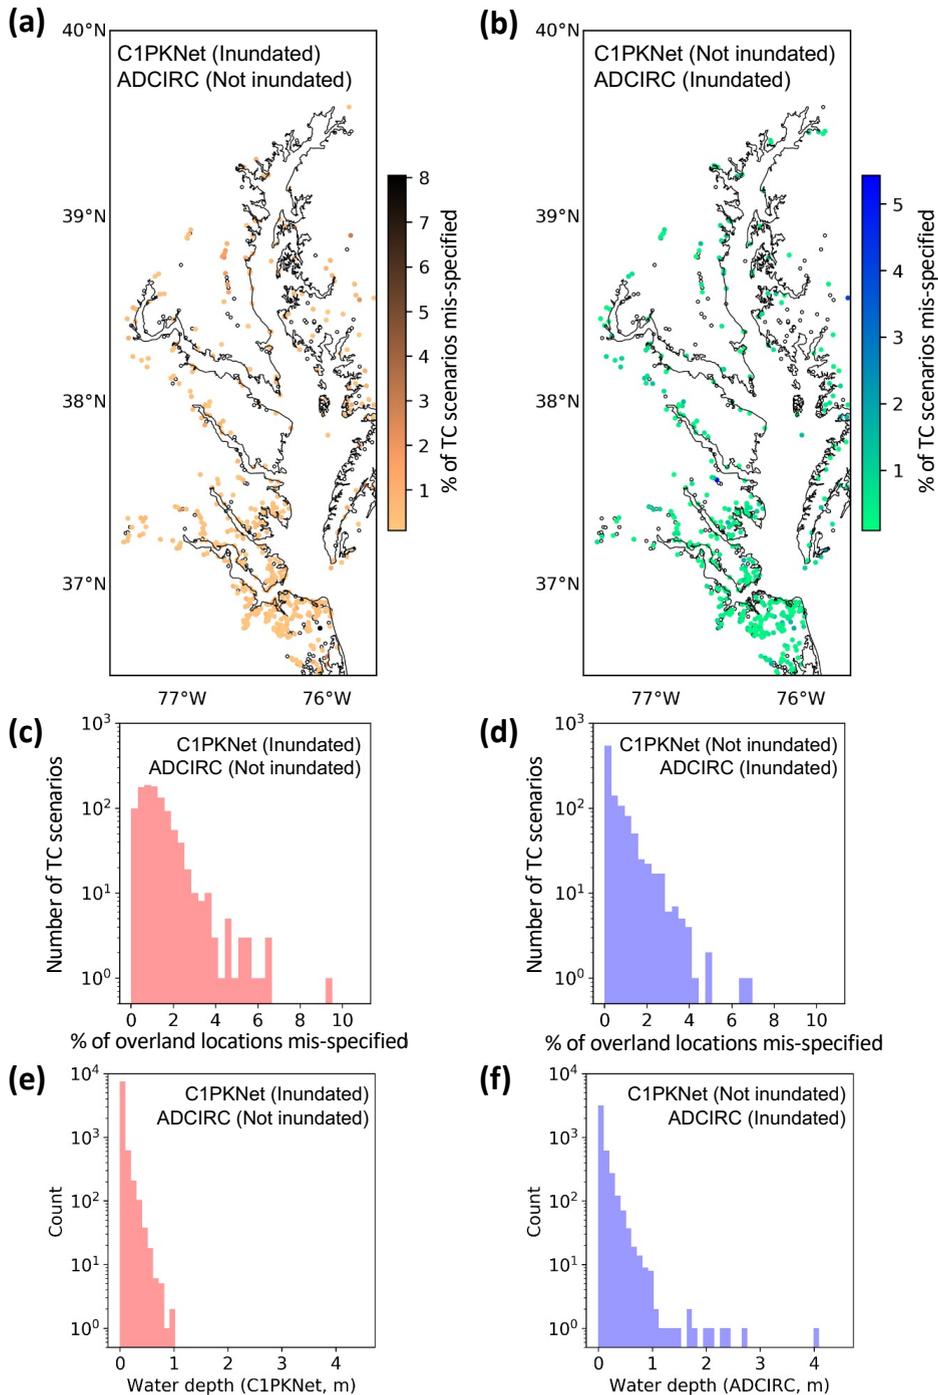
Figure 8: C1PKNet model performance in predicting the wet/dry condition at overland locations. (a) The percentage of TC scenarios, by overland location, where the C1PKNet model predicts flooding, but flooding is not predicted by ADCIRC simulation. Black circles represent zero. (b) The percentage of TC scenarios, by overland location, where the C1PKNet model does not predict flooding, but flooding is predicted by ADCIRC simulation. Black circles represent zero. (c,d) Histogram of the percentage of overland locations per scenario where the C1PKNet model predicts wet (c) or dry (d) conditions while ADCIRC simulations predict the opposite. (e,f) Histogram of the water depth (flood elevation less topographic elevation) based only on the data where the C1PKNet model predicts the wet (e) or dry (f) condition while ADCIRC simulations predict the opposite. The y-axis represents the number of data showing the corresponding water depth out of the dataset for 1031 synthetic TCs and the overland locations mis-specified as wet (e) or dry (f).
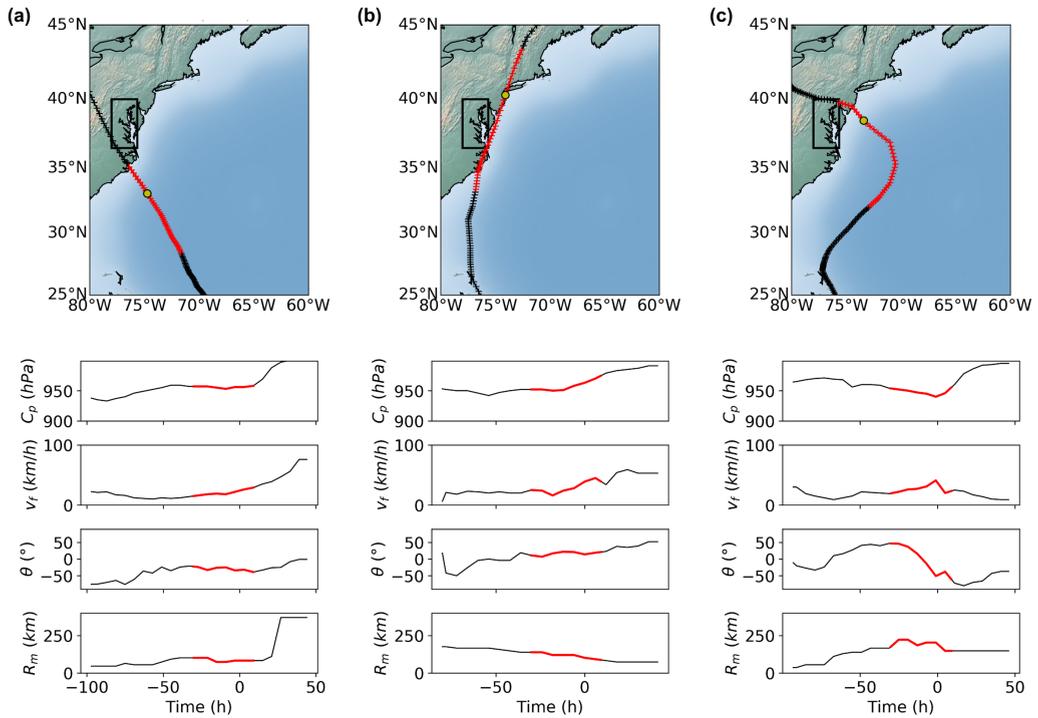
Figure 9: Time series of hurricane parameters for (a) Hurricane Isabel in 2003, (b) Hurricane Irene in 2011, and (c) Hurricane Sandy in 2012 (data from Demuth et al. (2006)). In the upper map, the yellow circle represents the reference point. The red lines are the 39 hours of hurricane parameters used as inputs in the C1PKNet model.
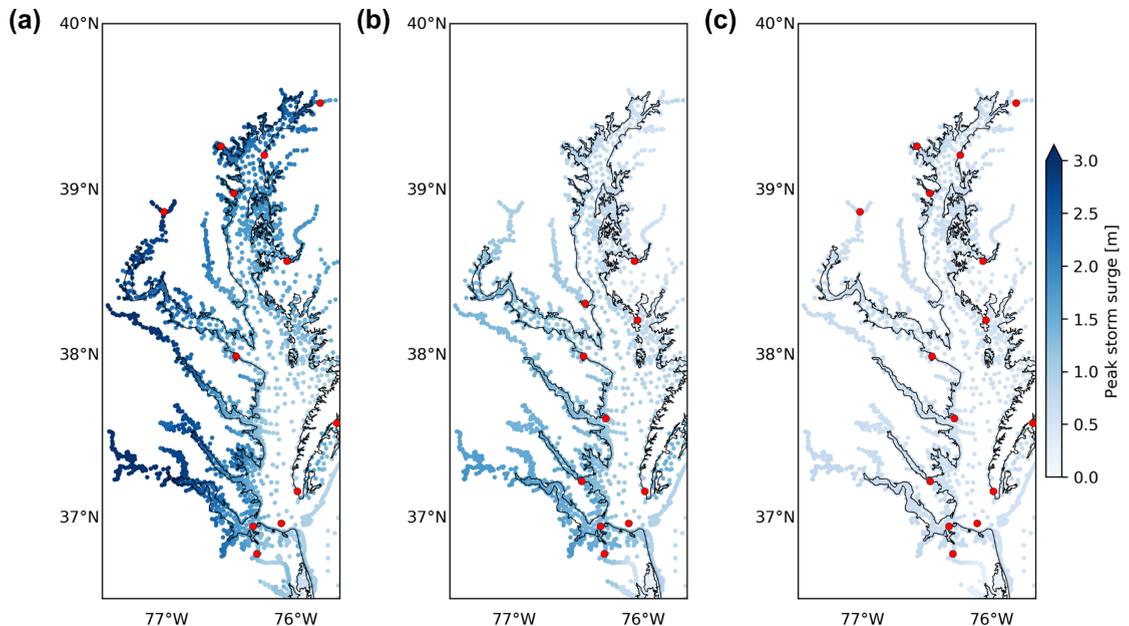
Figure 10: Peak storm surges predicted by the C1PKNet model for (a) Hurricane Isabel in 2003, (b) Hurricane Irene in 2011, and (c) Hurricane Sandy in 2012. The red circles represent the NOAA stations where the measured water-level data exists.
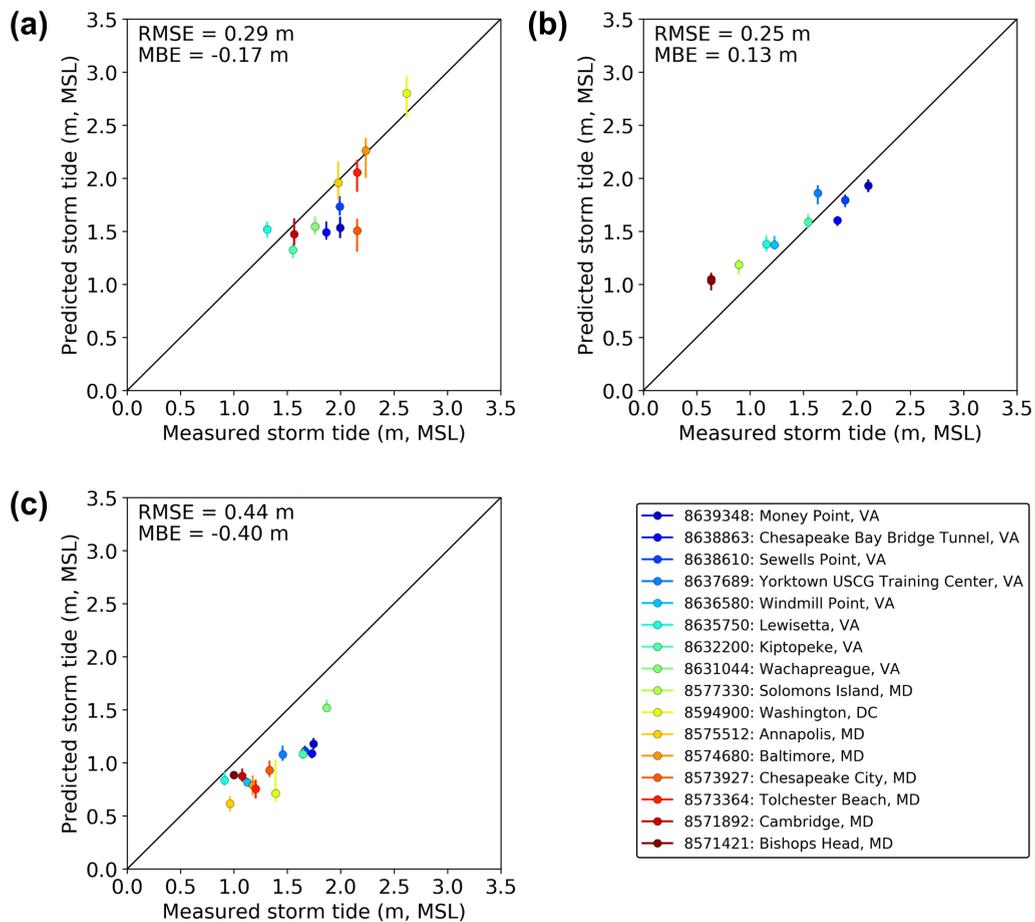
Figure 11: Comparison between measured storm tide (x-axis) and C1PKNet storm tide predictions (y-axis) for (a) Hurricane Isabel in 2003, (b) Hurricane Irene in 2011, and (c) Hurricane Sandy in 2012. In the y-axis, the circle and vertical line represent the mean and the range of 10 trained C1PKNet model predictions, respectively. The root mean square error (RMSE) and mean bias error (MBE) are denoted in each pane.
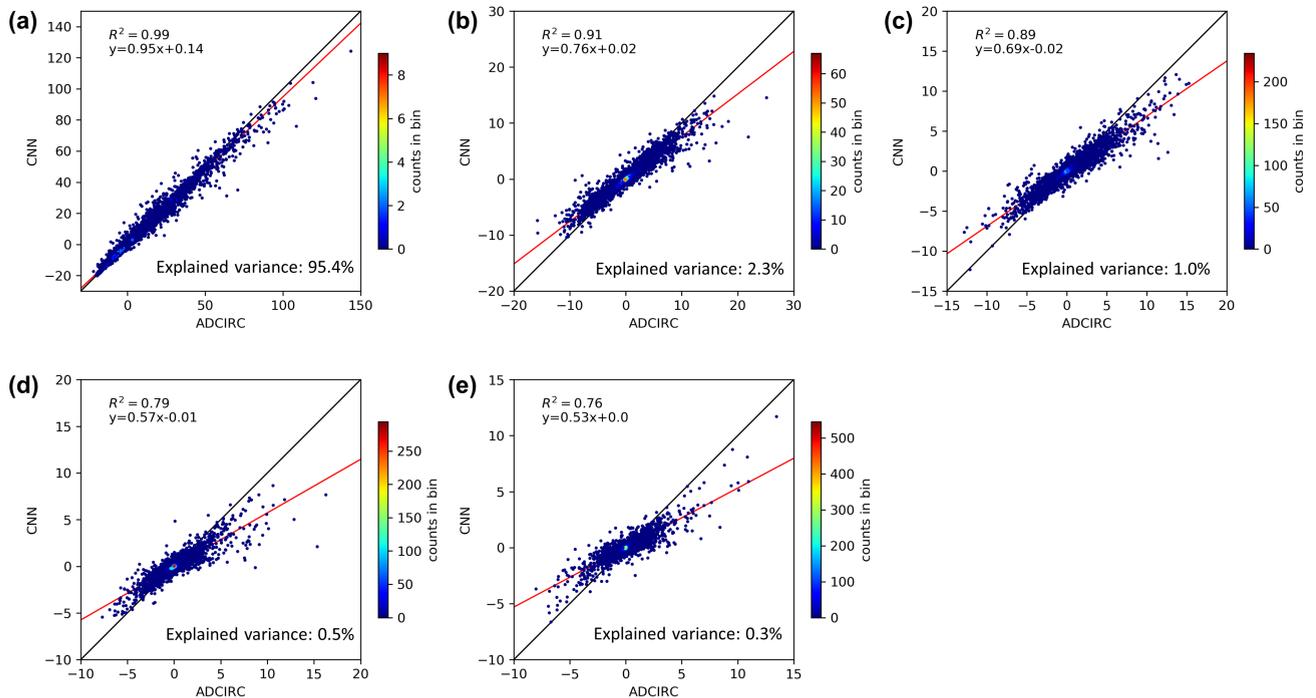
Figure 12: Performance of convolutional neural network (CNN). Principal components derived from the scaled peak storm surges of NACCS (x-axis) versus principal components predicted by the CNNs (y-axis) for (a) first principal component, (b) second principal component, (c) third principal component, (d) fourth principal component, and (e) fifth principal component. Each plot contains $12,372$ data values ($1,031$ storms × 12 clusters). The bin size of each axis is 0.05. The black line represents the equal line, and the red line is the linear regression solution. The best-fitting linear equation and coefficient of determination (R^2) are denoted in the upper left corner of each pane. The explained variance described by each component are shown in the lower right corner of each pane.
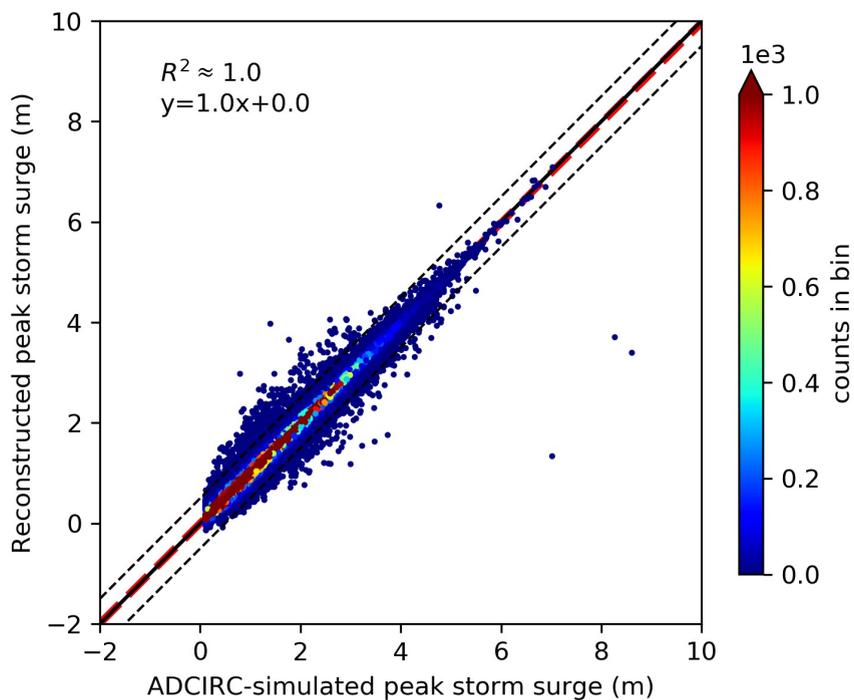
Figure 13: Comparison between the peak storm surges of NACCS (x-axis) and the peak storm surges reconstructed using five principal components (y-axis) for the first outer fold. The plot contains $2,883,897$ data values (927 storms $\times$ $3,111$ points). The bin size of each axis is 5 *cm*. The black line represents the equal line, the black dashed lines indicate ±0.5 *m* about an exact match, and the red dashed line is linear regression solution. The best-fitting linear equation and coefficient of determination (R^2) are denoted in the upper left corner.
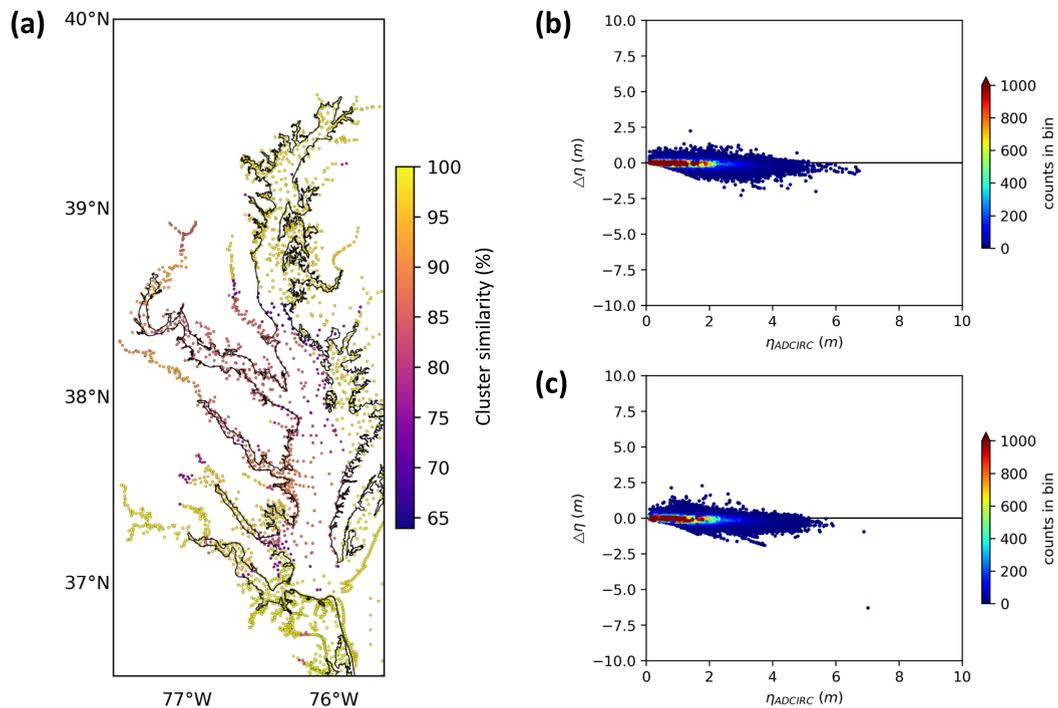
Figure 14: Performance of k-means clustering. (a) Map of cluster similarity (CS) that represents how often the location is classified into the same cluster. The larger the cluster similarity value, the more often the location is classified into the same cluster. (b and c) Density scatter plots showing the error distribution depending on the ADCIRC-simulated peak storm surge from the NACCS database. The bin size of each axis is 5 cm. The black horizontal line represents the zero error line. (b) Error distribution based on the locations where CS is larger than 98.1%. (c) Error distribution based on the locations where CS is smaller than 98.1%.