Regional surge hazard map development using Gaussian Process metamodeling

WoongHee Jung
Graduate Student, Department of Civil and Environmental Engineering and Earth Sciences, University of Notre Dame, Notre Dame, United States

Alexandros A. Taflanidis
Professor, Department of Civil and Environmental Engineering and Earth Sciences, University of Notre Dame, Notre Dame, United States

Norberto C. Nadal Caraballo
Senior Research Engineer, U.S Army Corps of Engineers, Research and Development Center Coastal Hazards Group, Vicksburg, United States

Madison C. Yawn
Research Engineer, U.S Army Corps of Engineers, Research and Development Center Coastal Hazards Group, Vicksburg, United States

Luke A. Aucoin
Research Engineer, U.S Army Corps of Engineers, Research and Development Center Coastal Hazards Group, Vicksburg, United States

ABSTRACT: The recent, very active hurricane seasons, as well as emerging concerns related to the future effects of sea level rise, storm intensification, and increased hurricane occurrence rate projections on coastal areas, make the prediction of storm-flood hazard a key priority when discussing coastal community resilience within planning (pre-disaster), emergency management, and post-disaster settings. To address this priority, researchers have placed substantial efforts in developing improved high-fidelity, numerical models to predict surges for a given storm. For promoting computational efficiency when utilizing these models within hazard estimation applications, surrogate modeling techniques have emerged as a popular strategy. The accuracy of surrogate modeling techniques in this context has been examined so far using cross-validation (CV) or test-sample validation techniques or by testing their performance for a (very) small number of historic storms. This paper investigates this topic within a different setting, examining the resultant regional storm surge hazard maps, specifically using Gaussian Process (GP) as metamodeling technique. This is accomplished by examining the hazard products obtained by GP implementations, as well as hazard products established through alternative, simplified approaches. Examining this accuracy fills in an important knowledge gap and provides an answer to the question "what are the real benefits in hazard estimation by using surrogate models?", improving at the same time trustworthiness of the associated results within the context of the coastal infrastructure risk quantification. A computationally efficient framework is also presented to explicitly consider the uncertainty associated with the GP predictions to provide confidence bounds for the hazard products.

1. INTRODUCTION

In hurricane-prone regions, tropical storm-induced surge represents one of the most significant flood components, especially when discussing hazard products associated with large return periods (Resio et al. 2007). Other important components may include, depending on the region, the effects of tides, extra-tropical storms,

and rainfall (i.e., compound flooding). The standard practice for predicting inundation probabilities associated with tropical storms is the following. Based on climatological data, an ensemble of storms is generated, typically through some variant of optimal sampling based on the joint probability method [JPM-OS] (Toro et al. 2010). This results in a set of storms and corresponding annual rates of occurrence $\{\mathbf{x}^h, \lambda^h; h = 1,..., N\}$, with $\mathbf{x}$ representing here an idealized representation of each storm, and λ denoting the rates. The storm surge for each storm is estimated through an appropriate hydrodynamic model, providing for the *i*th location in the region (typically corresponding to a node of that model) the surge estimate $z_i(\mathbf{x})$, expressed herein as a function of the parametric storm description. The annual rate of exceeding a threshold *b* can be then obtained utilizing the aforementioned storm ensemble as (Resio et al. 2007):

$$\begin{aligned} \lambda_i(b) &= 1 - \exp\left(-\sum_{h=1}^{N} \lambda^h \mathbb{I}[z_i(\mathbf{x}^h) > b]\right) \\ &\approx \sum_{h=1}^{N} \lambda^h \mathbb{I}[z_i(\mathbf{x}^h) > b] \end{aligned} \quad (1)$$

where $\mathbb{I}[.]$ is the indicator function, which is one if the quantity inside the bracket holds and zero if it does not, and the approximation in Eq. (1) holds for small values (<0.05) of $\lambda_i(b)$. Utilizing this prediction, the surge threshold b_i^p corresponding to different 1/*p* return periods of interest can be calculated so that $\lambda_i(b_i^p) = p$.

To support higher-accuracy hazard estimation, researchers have exerted significant efforts in the last decade in developing improved high-fidelity, numerical models to predict surge for a given storm (Luettich et al. 1992). Unfortunately, the computational burden associated with these models is significant, requiring a few thousand CPU hours for each simulation run, creating a very large computational cost for performing the ensemble analysis needed for the estimation represented by Eq. (1). This has motivated researchers to investigate the use of data-driven surrogate modeling techniques (Irish et al. 2009; Jia and Taflanidis 2013; Taflanidis et al. 2013; Lee et al. 2021), developed using a (small) database of high-fidelity storm simulations. Once calibrated, these surrogate models can replace the original numerical model to provide an inexpensive emulator for the storm surge. The hazard estimation can be subsequently performed with high computational efficiency using the surrogate model.

The accuracy of surrogate modeling techniques in this context has been traditionally examined by comparison of the predicted water levels for historical storms or through cross-validation (CV) techniques (Irish et al. 2009; Jia and Taflanidis 2013). This paper examines the accuracy from the perspective of the established hazard products, using Gaussian Process (GP) regression as surrogate modeling technique, since GPs have been shown in a number of studies (Jia et al. 2016; Zhang et al. 2018; Kyprioti et al. 2021; Kyprioti et al. 2022) to offer high-accuracy emulation for the storm surge, and, furthermore, serve as the foundation for many of the coastal risk assessment products for many federal agencies like the U.S Army Corps of Engineers (Nadal-Caraballo et al. 2020).

2. GP EMULATION FOR STORM SURGE

Detailed description of the development of a GP for storm surge emulation is included in (Jia et al. 2016; Kyprioti et al. 2021). Here the basic steps are reviewed. Formulation assumes that a database exists consisting of *n* simulations of the high-fidelity hydrodynamic model. Each simulation provides the predicted storm surge across a large number of n_z locations, typically corresponding to the grid of the numerical model. We will denote by $\mathbf{z} \in \mathbb{R}^{n_z}$ the storm surge vector and by z_i the surge at the *i*th location.

Step 1: Database parameterization. First, a parameterization of the database is needed to define a vector of storm features that will serve as the emulator input. This ultimately provides a n_x dimensional parametric description $\mathbf{x} \in \mathbb{R}^{n_x}$ for each storm. These features typically pertain to track, size, and intensity characteristics, and to account for the evolution of these characteristics over time, they may be (i) defined based on a specific reference instance (landfall or reference landfall for bypassing storms) (Jia et al. 2016), (ii)

averaged over some time-window around such a reference instance (Kyprioti et al. 2021), or (iii) described by considering the entire functional dependence over time (Lee et al. 2021). This provides the input matrix $\mathbf{X} \in \mathbb{R}^{N \times n_x}$ for the database, whose *h*th row corresponds to the parametric description $\mathbf{x}^h$ for the *h*th storm.

Step 2: Database imputation. For nearshore or onshore nodes that have remained dry in some of the storm simulations, the imputation of the database is performed. This imputation process provides the pseudo-surge, which then replaces the missing data in the original database. Exploiting the close proximity of nodes, it can be performed for each storm using spatial interpolation, for example using weighted *k* nearest neighbor (*kNN*) interpolation (Kyprioti et al. 2021). This provides the output matrix $\mathbf{Z} \in \mathbb{R}^{N \times n_z}$ for the database, whose *h*th row corresponds to *h*th storm pseudo-surge.

Step 3: Dimensionality reduction. To address the high dimensionality of the output **z**, Principal component analysis (PCA) is used as a dimensionality reduction technique (Jia and Taflanidis 2013). PCA identifies, through a linear projection, a small number of latent outputs (principal components) $\{y_k; k = 1,...,m\}$ that best explain the original observation matrix **Z**. For each component, it provides the projection vector $\mathbf{P}_k \in \mathbb{R}^{n_z}$ as well as the vector $\mathbf{Y}_k = [y_k(\mathbf{x}^1), \cdots, y_k(\mathbf{x}^N)]^T \in \mathbb{R}^N$ of observations over the storm database. The transformation also utilizes the mean of the observations for z_i, denoted herein as μ_i.

Step 4: Principal component GP calibration. For each of the principal components a separate GP is developed using input-output observation pair **X-Y**$_k$. The GP metamodel approximates the true response as a realization of a Gaussian process with correlation function between inputs $R_k(\mathbf{x}, \mathbf{x}')$, and mean trend expressed through a linear regression component. This leads to approximation $y_k \sim N(\tilde{y}_k(\mathbf{x} \mid \mathbf{X}), (\sigma_k^y(\mathbf{x} \mid \mathbf{X}))^2)$ where $N(a,b)$ stands for Gaussian with mean a and variance b, and $\tilde{y}_k(\mathbf{x} \mid \mathbf{X})$ and $(\sigma_k^y(\mathbf{x} \mid \mathbf{X}))^2$ correspond to the predictive mean and variance, respectively of the GP emulator. Details for these two functions as well as for the GP calibration, corresponding to the selection of the hyper-parameters of the correlation kernel $R_k(\mathbf{x}, \mathbf{x}')$, are included in (Kyprioti et al. 2021). Note that instead of individual surrogate models for each component, a grouping of the components may be considered as an alternative implementation to facilitate higher computational efficiency (Kyprioti et al. 2021).

Step 5: Storm surge predictions. Combining the predictions for all principal components, the surrogate model approximation for the storm surge is obtained using the linear PCA transformation. Adopting notation $[.]_{ik}$ to denote the $\{i,k\}$ element of a matrix, the mean and variance predictions for the *i*th node are:

$$\tilde{z}_i(\mathbf{x} \mid \mathbf{X}) = \mu_i + \sum_{k=1}^{m} [\mathbf{P}]_{ik} \tilde{y}_k(\mathbf{x} \mid \mathbf{X}) \quad (2)$$

$$\sigma_i^2(\mathbf{x} \mid \mathbf{X}) = \sum_{k=1}^{m} [\mathbf{P}]_{ik}^2 (\sigma_k^m(\mathbf{x} \mid \mathbf{X}))^2 \quad (3)$$

Since PCA is a linear transformation, the GP emulation establishes the probabilistic predictions $z_i \sim N(\tilde{z}_i(\mathbf{x} \mid \mathbf{X}), \sigma_i^2(\mathbf{x} \mid \mathbf{X}))$ (Jia and Taflanidis 2013). The mean of this distribution is commonly taken to represent the metamodel predictions, though for the hazard estimation, the use of the GP probabilistic description can be explored.

3. GP-ENABLED HAZARD ESTIMATION

3.1. Hazard estimation

The calibrated GP can be leveraged to support greater computational efficiency in the hazard estimation of Eq. (1). If only the mean GP predictions are utilized, i.e. when ignoring the GP variability, the approximation to the annual rate of exceedance threshold *b* is:

$$\tilde{\lambda}_i(b) \approx \sum_{h=1}^{N} \lambda^h \mathbb{I}[\tilde{z}_i(\mathbf{x}^h \mid \mathbf{X}) > b] \quad (4)$$

The variability in the GP predictions can be further utilized to estimate the expected value of this rate, $E_{GP}[\lambda_i(b)]$, where $E_{GP}[.]$ stands for expectation under the emulator probabilistic description. Using the total probability theorem, we can derive:

$$E_{GP}[\lambda_i(b)] \approx \sum_{h=1}^{N} \lambda^h \Phi\left(\frac{\tilde{z}_i(\mathbf{x}^h \mid \mathbf{X}) - b}{\sigma_i(\mathbf{x}^h \mid \mathbf{X})}\right) \quad (5)$$

where Φ(.) stands for the standard Gaussian cumulative distribution function. Note that for any storms for which the high-fidelity predictions are available, $\tilde{z}_i(\mathbf{x}^h \mid \mathbf{X})$ in Eqs. (4) and (5) should be replaced with $z_i(\mathbf{x}^h)$ while $\sigma_i(\mathbf{x}^h \mid \mathbf{X})$ in Eq. (5) should be set to 0.

Using either of these GP-based approximations for $\lambda_i(b)$ [i.e. either Eq. (4) or Eq. (5)], an approximation, denoted by $\tilde{b}_i^p$, for the threshold associated with different return periods $1/p$ can be obtained. The accuracy of this approximation can be estimated by calculating different error statistics over the nodes, for example, the correlation coefficient (dimensionless statistic) or surge score (dimensional statistic), given, respectively by:

$$CC(p) = \frac{\sum_{i=1}^{n_z}(b_i^p - \frac{1}{n_z}\sum_{i=1}^{n_z} b_i^p)(\tilde{b}_i^p - \frac{1}{n_z}\sum_{i=1}^{n_z}\tilde{b}_i^p)}{\sqrt{\sum_{i=1}^{n_z}(b_i^p - \frac{1}{n_z}\sum_{i=1}^{n_z} b_i^p)^2 \sum_{i=1}^{n_z}(\tilde{b}_i^p - \frac{1}{n_z}\sum_{i=1}^{n_z}\tilde{b}_i^p)^2}} \quad (6)$$

$$SC(p) = \frac{1}{n_z}\sum_{i=1}^{n_z} |\max(b_i^p, e_i) - \max(\tilde{b}_i^p, e_i)| \quad (7)$$

where e_i denotes the elevation of each node and max(a,b) corresponds to the maximum of the two arguments. Note that, as discussed in detail in (Kyprioti et al. 2021), the surge score incorporates the wet/dry node condition in the validation statistics.

3.2. Hazard confidence bounds

The GP probabilistic description can be additionally leveraged to obtain confidence bounds for the hazard curves for each node. This is established by creating realizations of the GP response across the storm ensemble used in the hazard estimation of Eq. (1), further considering the correlation dictated by the GP. This needs to be couched within the dimensionality reduction established through PCA. The process is the following. Let $\mathbf{u}_k = [y_k(\mathbf{x}^1), \cdots, y_k(\mathbf{x}^N)]^T$ denote the vector or responses for the kth principal component across the storm ensemble $\{\mathbf{x}^h; h = 1, ..., N\}$, with the hth element of $\mathbf{u}_k$ representing the response for the hth storm $y_k(\mathbf{x}^h)$. Based on the GP, $\mathbf{u}_k$ follows a Gaussian distribution $N(\tilde{\mathbf{u}}_k, \mathbf{C}_k)$. The hth element of mean vector $\tilde{\mathbf{u}}_k$ corresponds to the GP mean $\tilde{y}_k(\mathbf{x}^h \mid \mathbf{X})$, while the covariance matrix is assembled utilizing the GP variance and correlation. The elements of this matrix are:

$$[\mathbf{C}_k]_{hl} = \sigma_k^y(\mathbf{x}^h \mid \mathbf{X}) \cdot \sigma_k^y(\mathbf{x}^l \mid \mathbf{X}) \cdot R_k(\mathbf{x}^h, \mathbf{x}^l) \quad (8)$$

Let $\mathbf{u}_k^s$ denote a sample realization drawn from $N(\tilde{\mathbf{u}}_k, \mathbf{C}_k)$, with superscript s used to denote the sth sample. A sample for the surge response across the storm ensemble can be obtained by combining the samples from the different principal components as:

$$\begin{bmatrix} \tilde{z}_i^s(\mathbf{x}^1) \\ \dots \\ \tilde{z}_i^s(\mathbf{x}^N) \end{bmatrix} = \mu_i + \sum_{k=1}^{m} [\mathbf{P}]_{ik} \mathbf{u}_k^s \quad (9)$$

Utilizing this sample, a sample for the annual rate of exceeding threshold b is obtained as:

$$\lambda_i^s(b) \approx \sum_{h=1}^{N} \lambda^h \mathbb{I}[\tilde{z}_i^s(\mathbf{x}^h) > b] \quad (10)$$

Statistics for this rate, such as confidence bounds, can be obtained by generating a large number of samples s=1,…,N_c and calculating sample-based approximations for them. This implementation fully incorporates the probabilistic GP description in the estimation, including the correlation across the storm ensemble through the covariance matrix $\mathbf{C}_k$ for each principal component.

4. ILLUSTRATIVE CASE STUDY

4.1. Case study database

The database used in this study is part of the U.S. Army Corps of Engineers' Coastal Hazards System (CHS; https://chs.erdc.dren.mil). The CHS Louisiana Coastal Study (CHS-LA) was conducted for quantifying storm hazards and coastal compound flooding in Louisiana, including areas in the vicinity of the Greater New Orleans Hurricane Storm Damage Risk Reduction System (HSDRRS). The storm suite developed for CHS-LA consists of 645 synthetic tropical cyclones (TCs). Simulation of water levels for all TCs was performed using high-resolution, high-

fidelity atmospheric and hydrodynamic numerical models. The parameters of the synthetic TC suite were first used as input to drive a Planetary Boundary Layer (PBL) model on a nested grid to generate the wind and pressure fields used as forcing in the hydrodynamic modeling. The hydrodynamic simulations were performed by coupling the ADCIRC [ADvanced CIRCulation model] and the SWAN [Simulating Waves Nearshore] wave model. Further details for the database characteristics are included in (Kyprioti et al. 2022). For the storm parameterization, defining **x**, the following characteristics at landfall are used: the latitude and longitude for the storm landfall, the heading direction for the storm track, the central pressure deficit, the radius of maximum winds, and the translational speed. For the GP and subsequent hazard map development, the utilized grid consists of n_z=1,179,179 nodes while for the surge imputation a weighted *kNN* interpolation was used. Details for the GP development are included in (Kyprioti et al. 2022).

4.2. Case study implementation details

The accuracy of the hazard products obtained through the GP implementation is examined by considering the following setup. The reference hazard estimates b_i^p are established according to Eq. (1) using all available high-fidelity simulations, with N=645. A hypothetical scenario is then considered, in which the available computational resources were not sufficient for performing simulations for the entire ensemble of N storms, rather simulations can be for performed for only n storms (different values of n are examined ranging from 25 to 645). The terminology hazard-ensemble and reduced-ensemble will be utilized herein to describe this two sets. Utilizing the high-fidelity results for the n storms in the reduced-ensemble a GP is developed, and this GP is then used to provide predictions for the remaining storms within the hazard-ensemble and, finally, obtain the hazard estimates according to Eq. (4) or Eq. (5). These estimates will be referenced herein as GP-D (deterministic) and GP-P (probabilistic), respectively. This implementation evaluates hazard utilizing the entire hazard-ensemble but relies on high-fidelity simulations only for the reduced-ensemble, greatly reducing computational burden. The reduced-ensemble is chosen based on space filling principles within the original hazard-ensemble, utilizing k-means clustering for selecting the storms informing, ultimately, the GP development. To compare the accuracy of the GP-enabled hazard estimation, an alternative formulation is considered using strictly the high-fidelity simulations of the reduced-storm ensemble. In this case, hazard is given by Eq. (1) but using only the n available high-fidelity simulations for the reduced-ensemble. To accommodate this implementation a redistribution of the storm rates is necessary: the storm rates λ^h of the (N-n) storms that are removed from the hazard-ensemble, are attributed to the$n$ retained storms according to their proximity to them. The hazard estimates obtained through this process will be referenced herein as MC-R (reduced Monte Carlo).

4.3. Results and discussion

Figures 1 and 2 present accuracy statistic for all GP implementations and for MC-R for 4 different return periods 1/p=50,100,200,500 years for different values of n. Figure 1 presents results for the correlation coefficient and Figure 2 for the surge score. Figure 3 presents the surge maps for p=1/100 for the reference hazard estimation for GP-P, and MC-R using n=200. Results for GP-D are not presented here due to their similarity to GP-P. In Figure 3, the hazard maps established by the three different implementations as well as the relative errors between the approximate approaches (GP-P and MC-R) and the reference hazard estimate are separately provided across the two figure columns. Figures 1 and 2 present the average estimation accuracy over the entire domain, while Figure 3 presents the distribution of the estimation error across the different geographical locations.

Results in these figures clearly demonstrate the advantages offered by the GP-enabled hazard estimation framework, since for the same computational effort (same number of simulations n) it provides significant improvements in accuracy over the alternative implementation

(MC-R). As the number of storms in the ensemble increases (x-axis in Figures 1 and 2), the gap between MC-R and the GP-variants decreases, demonstrating that the latter approached (GP-based) will be particularly useful when the computational constraints are large, i.e. only when a small number of simulations can be performed. As expected the accuracy for all approximate approaches is higher for lower return periods (compare across the subplots in Figures 1 and 2) since they represent higher frequency events that can be estimated with higher confidence even for lower n values. But even for the higher return periods a value of n=150 to 200 seems sufficient for providing estimates with adequate accuracy (correlation coefficient over 95%), demonstrating the potential for a tremendous reduction in computational burden offered by the GP-based framework for hazard estimation (3-5 times smaller number of numerical simulations required).

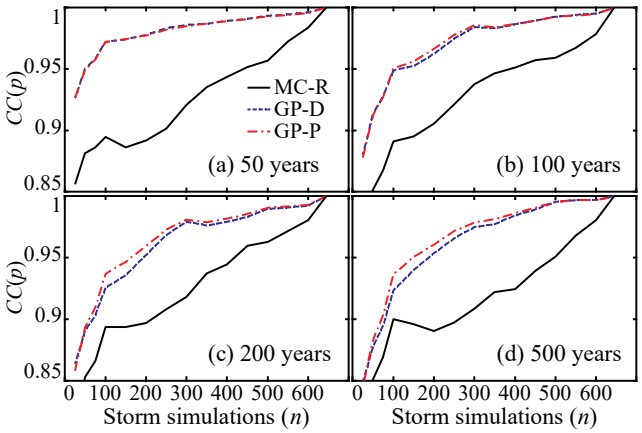

Figure 1: Correlation coefficient of different hazard approximation implementations against the number of storms available in the reduced-ensemble dataset. Results presented for four different return periods.

Comparing across the two different GP implementations in Figure 1 and 2, similar benefits are showcased, though for smaller number of simulations n the probabilistic estimation (GP-P) offers some small improvements in overall accuracy compared to the deterministic estimation (GP-D). This demonstrates a degree of robustness for GP-P when the number of available high-fidelity simulations is not sufficient for accommodating improved accuracy estimates. For values of n that one would consider in practice, values that offer hazard estimates with higher accuracy (for example, correlation coefficient of over 95%) the discrepancies between the two different GP implementations are small.

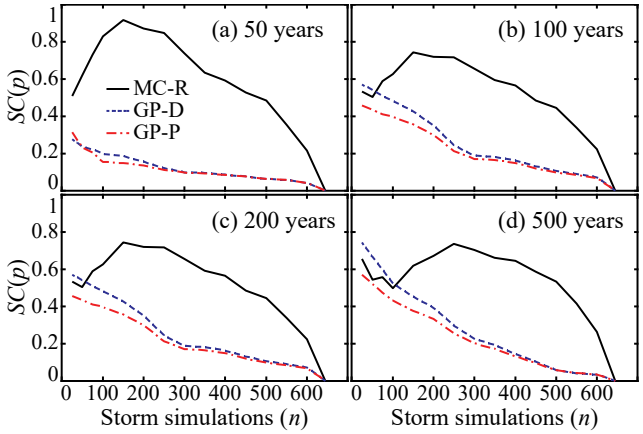

Figure 2: Surge score of different hazard approximation implementations against the number of storms available in the reduced-ensemble dataset. Results presented for four different return periods.

The spatial distribution of errors in Figure 3 reveals some very interesting additional trends. The GP-based estimation established low magnitude errors across the entire geographic domain, with absolute relative error less than 0.5 m across all locations even for implementation using only n=200 storm simulations. The same does not hold for the alternative implementation (MC-R) that not only produces larger errors (as also evident in Figures 1 and 2), but demonstrates a substantial higher variability of these errors within the geographic domain, with certain regions exhibiting higher relative errors. This is a very important point of comparisons as the GP-based implementations provide reliable estimates over the entire domain: there do not seem to be any subdomains that the hazard estimates consistently demonstrate greater errors. This means that the GP formulation targets a global improvement in accuracy, perhaps infused through the dimensionality reduction PCA step,

and does not indirectly prioritize improvements in certain geographical domains while sacrificing accuracy in others (to accomplish the aforementioned improvements). The MC-R implementation seems to indirectly establish such a compromise: some domains suffer from disproportionally reduced accuracy.

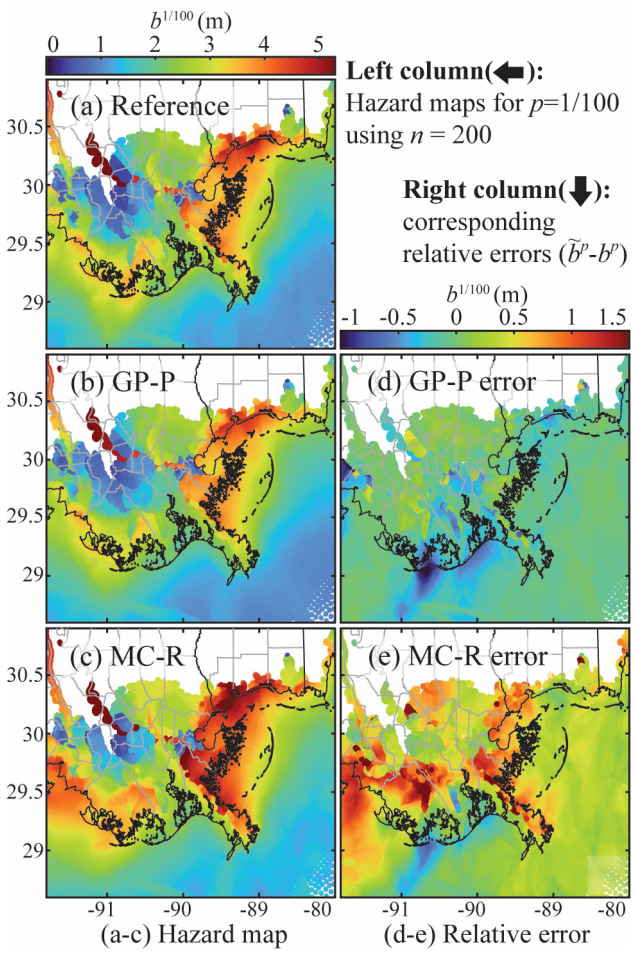

Figure 3: [left column] Hazard maps for 100 year return period established by (a) reference estimation and (b) GP-P and (c) MC-R using n=200 simulations. [right column] Relative error of the hazard approximation implementations.

Finally, Figure 4 presents the hazard curves $\lambda_i(b)$, including the confidence bounds established through the implementation discussed in Section 3.2, for four specific nodes. For obtaining the confidence bounds in this figure, a large number of N_c=10,000 samples were used, whereas GP implementation corresponds to n=200 number of simulations.

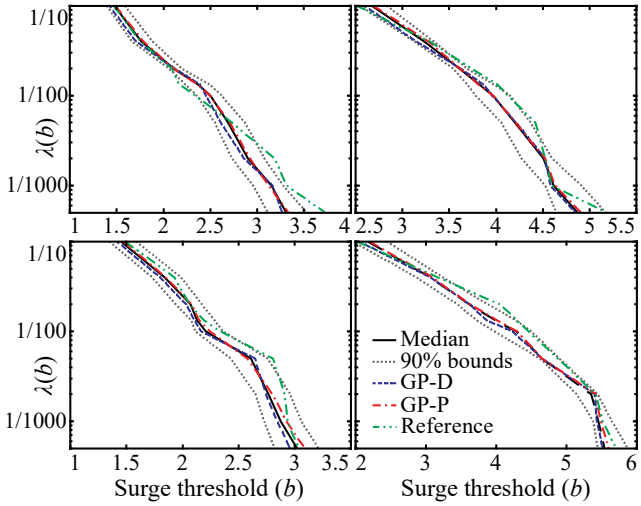

Figure 4: Hazard curves and their 90% confidence bounds established using N_c=10,000 realizations of the GP responses for four different nodes. GP implementation corresponds to n=200 storm simulations.

Results in Figure 4 demonstrate how the probabilistic characteristics of the GP facilitate a more comprehensive estimation of the hazard, providing confidence bounds that reflect the degree of uncertainty stemming from the GP approximation. In most instances the reference hazard estimate falls well within the provided confidence bounds. A deviation from this trend might exist for rare events (larger return periods), something that is not entirely unexpected due to the challenges, as discussed earlier, in accurately approximating such events. Nevertheless, the ability to provide realistic confidence bounds is significant advantage of the GP-based hazard approximation, since this is offered with moderate only increase in the computational complexity.

Finally, an interesting trend to note is that the GP-P predictions (corresponding to the mean exceedance rates as discussed earlier) and the median of the hazard curves do not match in Figure 4. This is not unexpected and is attributed to the correlation within the GP, incorporated in the estimation through function $R_k(\mathbf{x},\mathbf{x}')$ in the description of the covariance matrix across the storm ensemble in Eq. (8). This correlation across the considered storm ensemble will create differences between the mean and median estimates.

5. CONCLUSIONS

The implementation of Gaussian Process (GP) surrogate modeling for estimation of storm surge hazard maps was discussed in this paper. The accuracy of GPs in this context has been traditionally examined by comparison of the predicted water levels for historical storms or through cross-validation techniques. Here this validation was examined from the perspective of the established hazard products. Hazard estimates obtained by a GP implementation for different number of available storm simulations (informing the GP development) was compared to reference values. An alternative estimation using only the available storm simulations was also considered. Results in the case study clearly demonstrate the advantages offered by the GP-enabled hazard estimation framework, since for the same computational effort it provides significant improvements in accuracy over the alternative implementation. Even for larger return periods (infrequent events), a value of 150 to 200 storm simulations seems sufficient for providing estimates with adequate accuracy. A computationally efficient framework was also discussed to explicitly consider the uncertainty associated with the GP predictions to provide confidence bounds for the hazard products. It was shown that this facilitates a more comprehensive estimation of the hazard, providing confidence bounds that reflect the degree of uncertainty stemming from the GP approximation.

6. ACKNOWLEDGMENTS

This research was funded by the US Army Corps of Engineers (USACE), under the grant number W912HZ-22-C-0041.

7. REFERENCES

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(1), 183-205.

Jia, G. and 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, 24-38.

Jia, G., Taflanidis, A.A., Nadal-Caraballo, N.C., Melby, J.A., Kennedy, A.B., and Smith, J.M. (2016). "Surrogate modeling for peak and time dependent storm surge prediction over an extended coastal region using an existing database of synthetic storms", *Natural Hazards*, 81(2), 909-938.

Kyprioti, A.P., Taflanidis, A.A., Nadal-Caraballo, N.C., Yawn, M.C., and Aucoin, L.A. (2022). "Integration of node classification in storm surge surrogate modeling", *Journal of Marine Science and Engineering*, 10(4), 551.

Kyprioti, A.P., Taflanidis, A.A., Plumlee, M., Asher, T.G., Spiller, E., Luettich, R.A., Blanton, B., Kijewski-Correa, T.L., Kennedy, A.B., and Schmied, L. (2021). "Improvements in storm surge surrogate modeling for synthetic storm parameterization, node condition classification and implementation to small size databases", *Natural Hazards*, 109, 1349-1386.

Lee, J.W., Irish, J.L., Bensi, M.T., and Marcy, D.C. (2021). "Rapid prediction of peak storm surge from tropical cyclone track time series using machine learning", *Coastal Engineering*, 170, 104024.

Luettich, R.A., Westerink, J.J., and Scheffner, N.W. (1992). *ADCIRC: An advanced three-dimensional circulation model for shelves, coasts, and estuaries. Report 1. Theory and methodology of ADCIRC-2DDI and ADCIRC-3DL.* Dredging Research Program Technical Report DRP-92-6, U.S Army Engineers Waterways Experiment Station, Vicksburg, MS.

Nadal-Caraballo, N.C., Campbell, M.O., Gonzalez, V.M., Torres, M.J., Melby, J.A., and Taflanidis, A.A. (2020). "Coastal Hazards System: A Probabilistic Coastal Hazard Analysis Framework", *Journal of Coastal Research*, 95(SI):1211-1216.

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., Suhada, J.N., Toro, G.R., and Vickery, P.J. (2007). *White paper on estimating hurricane inundation probabilities*, Consulting Report prepared by USACE for FEMA.

Taflanidis, A.A., Kennedy, A.B., Westerink, J.J., Smith, J., Cheung, K.F., Hope, M., and Tanaka, S. (2013). "Rapid assessment of wave and surge risk during landfalling hurricanes; probabilistic approach", *ASCE Journal of Waterway, Port, Coastal and Ocean Engineering*, 139(3), 171-182.

Toro, G.R., Resio, D.T., Divoky, D., Niedoroda, A.W., and Reed, C. (2010). "Efficient joint-probability methods for hurricane surge frequency analysis", *Ocean Engineering*, 37(1), 125-134.

Zhang, J., Taflanidis, A.A., Nadal-Caraballo, N.C., Melby, J.A., and Diop, F. (2018). "Advances in surrogate modeling for storm surge prediction: storm selection and addressing characteristics related to climate change", *Natural Hazards*, 94(3), 1225-1253.