14th International Conference on Structural Safety and Reliability - ICOSSAR'25
1-6 June, 2025 - Los Angeles California, USA
R. Ghanem, E. Johnson, S. Masri and A. Olivier (Eds)

ON THE CHALLENGES OF ADOPTING ADAPTIVE MONTE CARLO TECHNIQUES IN REGIONAL RISK ASSESSMENT

WOONGHEE JUNG[1] AND ALEXANDROS A. TAFLANIDIS[2]

[1]Department of Civil and Environmental Engineering and Earth Sciences, University of Notre Dame
Notre Dame, IN 46556
wjung2@nd.edu

[2]Department of Civil and Environmental Engineering and Earth Sciences, University of Notre Dame
Notre Dame, IN 46556
a.taflanidis@nd.edu

Key words: adaptive importance sampling, adaptive multi-fidelity Monte Carlo, probabilistic storm surge estimation, high-dimensional output.

Abstract. This paper discusses the challenges encountered when adopting adaptive Monte Carlo (MC) techniques in regional risk assessment settings. The MC estimation needs to be established in this case for a very high-dimensional output, corresponding to different locations/assets within the region. The challenges originate from the fact that for the adaptive characteristics of the MC algorithms, a compromising solution needs to be established across the quantities of interest (QoIs) representing these outputs. Two different algorithms are discussed, adaptive importance sampling (AIS) and adaptive multi-fidelity Monte Carlo (AMFMC), for a specific application, the real-time probabilistic predictions for the anticipated surge for landfalling tropical storms/cyclones. These probabilistic predictions are made through an uncertainty quantification process that involves: (i) generating a sufficiently large ensemble of storm scenarios based on the nominal storm advisory and the anticipated forecast errors; (ii) performing high-fidelity numerical simulations to obtain surge predictions for each storm scenario; and (iii) estimating surge statistics of interest by assembling the simulation results. This process is repeated whenever the nominal storm advisory is updated. AIS and AMFMC are considered here to improve the uncertainty propagation efficiency, and reduce number of storm scenarios utilized in the analysis. A direct comparison between these frameworks is established, focusing on the challenges encountered in tuning the algorithm adaptive characteristics to provide probabilistic estimates across a large number of QoIs, corresponding to the surge predictions for different locations within the coastal region of interest. Different compromise solutions are promoted to accommodate the conflicting preferences the QoIs represent. The efficacy of the two frameworks is examined in detail in this setting, comparing the accuracy of idealized implementations (adaptive decisions independently made for each QoI) to the accuracy of practical implementations (single, compromise decision within the MC implementation). The study also showcases the importance of information sharing across storm advisories in real-time probabilistic storm surge predictions and provides guidelines for an efficient adaptive MC formulation in such settings.

1 INTRODUCTION

Advancements in computational modeling and probabilistic analysis techniques are enabling the assessment of natural hazards and their impact on the built environment and communities with unprecedented scale and resolution[1]. The outcome of such analyses becomes an invaluable ingredient in guiding emergency responses, assessing societal consequences, simulating the recovery phase, and optimizing disaster policy and design decisions. At the same time, such an attempt at large-scale regional risk assessment presents opportunities for researchers to tackle previously unexplored challenges in the field of uncertainty quantification (UQ). This paper examines one of these challenges, related to the adoption of adaptive Monte Carlo (MC) techniques in regional risk assessment. The specific MC algorithms discussed are: (i) adaptive importance sampling (AIS)[2]; and (ii) adaptive multi-fidelity Monte Carlo (AMFMC)[3]. The adaptive characteristics pertain, respectively, to the selections of: (i) the optimal importance sampling (IS) density; or (ii) the optimal budget allocation between the different-fidelity models. The challenges originate from the fact that for these characteristics, a compromising solution needs to be established across the quantities of interest (QoIs) corresponding to outputs for different locations/assets within the region. The number of such QoIs in regional risk assessment settings[1, 4] can easily exceed hundreds of thousands, creating a high-dimensional output problem that is not typically encountered beyond such settings[5].

Though the methodological developments discussed are general, they are couched within a specific application, the assessment of storm surge risk during landfalling events[6]. Accurate assessment of this risk is recognized as a critical component for enhancing the resilience of coastal communities due to its vital role in informing emergency response decisions[7]. This assessment is established by combining the national hurricane center (NHC) storm advisories with a high-resolution hydrodynamic model, like ADCIRC[8], that can provide accurate estimates of the anticipated surge using as input a parametric wind field based on these advisories. To guide robust decision-making, potential forecast errors in the storm advisories need to be explicitly considered[6]. This results in probabilistic predictions of the storm surge impacts, established in a MC setting using an ensemble of candidate storm scenarios[9]. This process is repeated whenever the nominal storm advisory is updated, typically every 6 hours. The number of storm scenarios utilized in the MC estimation impacts the statistical accuracy of the probabilistic predictions. However, the significant computational burden of the ADCIRC high-fidelity simulations and the necessity to deliver expedited real-time predictions limit this number, constraining the statistical accuracy. The adoption of the aforementioned adaptive MC techniques, AIS and AMFMC is examined here to address this challenge. Both are deployed by leveraging readily available high-fidelity simulation results from the probabilistic estimation performed for past advisories, to achieve the desired statistical accuracy improvement for the current advisory. AIS[10] utilizes simulation results from the previous advisory to adaptively formulate the IS density to generate a more informative storm ensemble with reduced size. AMFMC[11] adaptively develops a surrogate model, serving as a low-fidelity approximation within the multi-fidelity setting, using simulation results from multiple previous advisories.

This paper examines the challenges in tuning the adaptive characteristics for the two adaptive MC frameworks to support surge predictions across high-dimensional QoIs, corresponding to the surge impacts for different locations within the coastal domain. The importance of information sharing across storm advisories to accomplish the desired MC adaptivity is also showcased.

2 OVERVIEW OF REAL-TIME PROBABILISTIC SURGE HAZARD PREDICTION

As a storm approaches landfall, the NHC issues advisories that provide its forecasted features[6]: (i) track, described by the latitude, s_{lat}, and longitude, s_{lon}; (ii) size, described by the radius of maximum winds, R_{mw}; (iii) intensity, described by the maximum sustained wind speed, v_w. To formalize the UQ framework, define the four-dimensional time-dependent vector of storm features as $\mathbf{q}(t)=[R_{mw}(t)\ s_{lat}(t)\ s_{lon}(t)\ v_w(t)]^T$. For the ($k$)th advisory, let $\overline{\mathbf{q}}^{(k)}(t)$ denote the nominal forecasts for these storm features with the superscript $\cdot^{(k)}$ representing the advisory number. The uncertainty of these forecasts, denoted as $\mathbf{\Delta q}(t)$, is probabilistically characterized through statistical analysis of historical forecast errors[6], with some assumptions for the correlation of the temporal evolution of the storm features. This leads to a time-independent random variable vector $\mathbf{x}$, describing the deviation of future storm features from their nominal predictions, with its probability model $p(\mathbf{x})$ established based on the aforementioned statistical analysis[9]. The vector $\mathbf{x}$ defines $\mathbf{\Delta q}(t)$, while the combination of $\mathbf{\Delta q}(t)$ and $\overline{\mathbf{q}}^{(k)}(t)$ leads to a sample storm scenario, $\mathbf{q}(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$, for the ($k$)th advisory.

Each storm scenario can be translated into parametric input for a storm surge numerical model (like ADCIRC) to provide predictions for the anticipated surge across n_z discretized locations within the geographic domain where the storm is expected to impact. Let $z_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ ($i=1,...,n_z$) denote the predicted peak surge level at the ith location. To support emergency response decisions, different statistics of interest for the peak storm surge can be estimated[6], with the most popular ones being the expected (mean) surge or the probability that the surge will exceed a specific threshold b. To generalize the discussion, let $h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ denote the consequence measure for the ith location corresponding to the statistical product of interest. The latter is quantified by the probabilistic integral:

$$H_i^{(k)}=E_p[h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]=\int h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))\,p(\mathbf{x})d\mathbf{x} \tag{1}$$

where $E_p[.]$ represents the expectation operator under $p(\mathbf{x})$, $h_i(.)$ is defined according to the estimated statistical product, for example, $z_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ for the mean response or $I[z_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))>b]$ for the probability $P_i(b)$ that the surge will exceed a specific threshold b, where $I[\cdot]$ is the indicator function. The other common product is the surge threshold $b_i^{p_t}$ for a specific exceedance probability p_t. This is easily estimated by the inverse problem, $P_i(b_i^{p_t})=p_t$.

For the MC estimation of the probabilistic integral in Eq. (1), a sample set $\{\mathbf{x}^l: l=1,\ldots,N\}\sim p(\mathbf{x})$ is generated, leading to a storm ensemble $\{\mathbf{q}(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t)):l=1,...,N\}$. The surge response across all locations is then estimated for each storm, providing the corresponding consequence measures. The MC estimate of Eq. (1) and its variance (representing statistical accuracy) are[12]:

$$\hat{H}_i^{(k)}=(1/N)\sum_{l=1}^{N}h_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t)) \tag{2}$$

$$Var[\hat{H}_i^{(k)}]=Var_p[h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]/N=(E_p[h_i^2(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]-(H_i^{(k)})^2)/N \tag{3}$$

where $Var[.]$ denotes the variance operator and $Var_p[.]$ denotes the respective operator under $p(\mathbf{x})$. Note that each $h_i(.)$ represents a different QoI within this UQ setting, and that the estimation for all QoIs is simultaneously performed using the same numerical simulations. As discussed in the introduction, due to the significant computational burden of the numerical model, the value of N is typically kept small[6, 9], creating challenges in attaining accurate estimates. This incentivizes the introduction of the variance reduction techniques discussed next.

3 ADAPTIVE MC ESTIMATION THROUGH INFORMATION SHARING ACROSS ADVISORIES

3.1 AIS formulation and challenge in selecting adaptive characteristics

AIS introduces an IS proposal density, denoted as $f^{(k)}(\mathbf{x})$ for the (k)th advisory, to generate an ensemble with higher quality information for estimating the surge statistics of interest[12]. This density is adaptively chosen using simulation results from the previous advisory[10]. The replacement of $p(\mathbf{x})$ with $f^{(k)}(\mathbf{x})$ in the MC estimation leads to a sample set $\{\mathbf{x}^l: l=1,\ldots,N\}\sim f^{(k)}(\mathbf{x})$, and to the IS-based MC estimator for the ith QoI[12]:

$$\hat{H}_{i,IS}^{(k)} = (1/N)\sum_{l=1}^{N} h_i(\mathbf{x}^l \mid \overline{\mathbf{q}}^{(k)}(t))\frac{p(\mathbf{x}^l)}{f^{(k)}(\mathbf{x}^l)} \tag{4}$$

The IS estimator $\hat{H}_{i,IS}^{(k)}$ is guaranteed to be unbiased if the support of $f^{(k)}(\mathbf{x})$ is greater than that of the integrand $h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))p(\mathbf{x})$. Given this, the variance of the IS estimator is[12]:

$$Var[\hat{H}_{i,IS}^{(k)}] = \frac{1}{N}\left[E_{f^{(k)}}\left[\left(h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))\frac{p(\mathbf{x})}{f^{(k)}(\mathbf{x})}\right)^2\right] - (H_i^{(k)})^2\right] \tag{5}$$

where $E_{f^{(k)}}[\cdot]$ is the expectation operator under $f^{(k)}(\mathbf{x})$. The optimal IS density is the one minimizing this variance (maximizing statistical accuracy), given by[10, 12]:

$$f_i^*(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t)) \propto |\, h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))\,|\, p(\mathbf{x}) \tag{6}$$

The use of this optimal density in the MC estimation is impractical as it requires complete information for the surge responses at the (k)th advisory, which is unavailable. The AIS implementation in [10] proposes to leverage readily available simulation results from the previous (k–1)th advisory, $\{h_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k-1)}(t)): l=1,\ldots,N\}$, to establish an IS proposal density for the current advisory. This is achieved by using the results to introduce a sample-based approximation $\hat{f}_i^*(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k-1)}(t))$ of the optimal IS for the previous advisory, which is then utilized for IS proposal density selection for the current advisory[10]. The details for the implementation with an efficient approach through dimensionality reduction can be found in [10]. Note that the AIS formulation focuses its adaptive decisions on the random variable vector $\mathbf{x}$, while the information shared across advisories is the consequence measure, $h_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t))$. Due to its explicit dependence on the advisory $\overline{\mathbf{q}}^{(k)}(t)$, simulation results from only the previous advisory are used, as these provide the most relevant information for the current advisory.

As indicated by the dependence of the optimal IS density of Eq. (6) on the subscript $\cdot_i$, each QoI leads to a different IS density choice. This is demonstrated in Figure 1, depicting the distribution over the geographic domain of the mean of the optimal IS density for Superstorm Sandy advisory 26 for 10% exceedance probability statistics estimation, for all components of $\mathbf{x}$. Note that this specific advisory and storm is target of the case study discussed in Section 4. It is important to mention that the optimal IS density mean dictates the shift of the IS proposal density towards the peak of the probabilistic integrand[10]. Great variability is observed across the geographic domain, while QoIs promote conflicting decisions, shifting the original zero mean toward positive values for some QoIs and toward negative values for others.

Since a single storm ensemble needs to be utilized to facilitate the numerical estimation of all QoIs, a single IS proposal density needs to be selected, balancing across the conflicting QoI preferences. This compromise solution will reduce the effectiveness of the AIS implementation.

Study[10] offers extensive discussions on establishing the compromise solution and the promoted option is to choose the representative IS density, denoted by $\hat{f}_{IS}(\mathbf{x})$, as:

$$\hat{f}_{IS}(\mathbf{x}) = \sum_{i=1}^{n_z} \omega_i \hat{f}_i^*(\mathbf{x} \mid \bar{\mathbf{q}}^{(k-1)}(t)) / \sum_{i=1}^{n_z} \omega_i \tag{7}$$

where ω_i defines the relative importance (i.e., weight) of the ith QoI. It was recommended in [10] to utilize the variance of the output for the weight, i.e., $\omega_i = Var[h_i(\mathbf{x} \mid \bar{\mathbf{q}}^{(k-1)}(t))]$, as this gives higher importance to QoIs with greater statistical variability.

Finally, to guarantee an unbiased MC estimation through Eq. (4), a robustness issue related to the coverage of the representative IS density $\hat{f}_{IS}(\mathbf{x})$ has to be addressed. This issue originates from the fact that this density is selected using information from the previous (k–1)th advisory, which does not guarantee that its support will be greater than the integrand for the current (k)th advisory[10]. A defensive IS scheme[5] is adopted to address this challenge as:

$$\hat{f}_{DIS}(\mathbf{x} \mid \alpha) = \alpha \hat{f}_{IS}(\mathbf{x}) + (1-\alpha) p(\mathbf{x}) \tag{8}$$

with $0 \le \alpha \le 1$ representing the defensive IS weight. The use of $\hat{f}_{DIS}(\mathbf{x} \mid \alpha)$ guarantees unbiasedness, but reduces efficacy compared to the use of $\hat{f}_{IS}(\mathbf{x})$. An adaptive selection for α is established[10] by comparing the (reduced) projected efficiency of $\hat{f}_{DIS}(\mathbf{x} \mid \alpha)$ to the (original) projected efficiency of $\hat{f}_{IS}(\mathbf{x})$, and adopting an appropriate weight α_{DIS} so that the resulting reduction is less than a threshold γ. The final density $\hat{f}_{DIS}(\mathbf{x} \mid \alpha_{DIS})$ eventually serves as the IS proposal density. This density establishes not only a rational compromise across the QoIs but also promotes robustness with only a moderate sacrifice for the projected efficiency.

It is important to note that the IS proposal density changes with the surge statistic of interest (i.e., the definition of the consequence measure). Though the IS density selected for an event with moderate likelihood showed overall good performance across different statistical products in [10], it is ideal to use a tailored IS density to the surge statistic of interest.

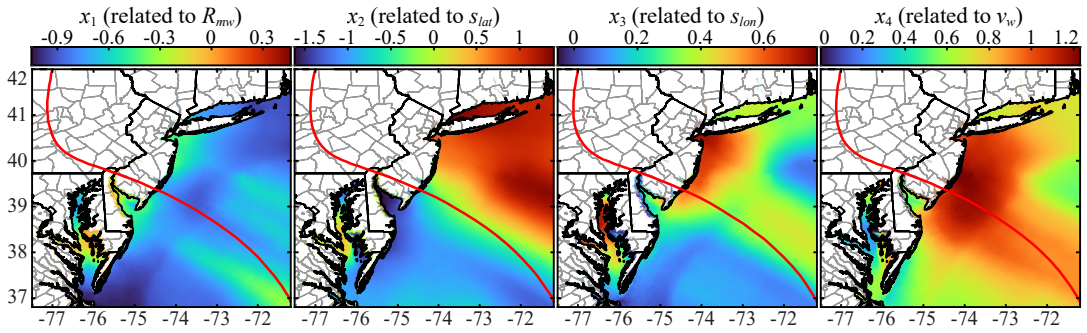

Figure 1: Spatial variability of the mean of the optimal IS density corresponding to 10% exceedance probability statistics for Superstorm Sandy advisory 26. Each column corresponds to a different component of vector **x**. The nominal storm track is also shown with a red line.

3.2 AMFMC formulation and challenge in selecting adaptive characteristics

AMFMC introduces a low-fidelity model to accelerate the MC estimation within MFMC setting[3]. The formulation in [11] uses as low-fidelity model a surrogate model approximation, which is adaptively developed based on high-fidelity simulation results from multiple previous

advisories. This surrogate model serves as a computationally efficient, data-driven approximation for the input/output relationship of the high-fidelity model. The shared simulation results correspond to the surge responses rather than the consequence measures appearing in the probabilistic integrals, as the estimation of the responses is the most computationally intensive part in calculating the surge statistics of interest. Since the surge response $z_i(\mathbf{q}(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t)))$ is not explicitly dependent on $\overline{\mathbf{q}}^{(k)}(t)$, AMFMC has greater flexibility compared to AIS, allowing information from multiple past advisories, instead of a single advisory, to be utilized. Additionally, since implementation requires a small number of high-fidelity simulations from the current advisory, the so-called pilot simulations, these simulations can also support the adaptive development of the surrogate model. To avoid bias when incorporating information from the current advisory[3], using leave-one-out (LOO) predictions for the current advisory storms is proposed in [11].

Once the surrogate model is developed, AMFMC utilizes surge predictions from both the low-fidelity and the high-fidelity models to obtain a MC estimate of the probabilistic integral in Eq. (1). The use of the low-fidelity simulations accelerates the MC estimation (improves statistical accuracy) by exploiting the correlation between the two models, while the use of the high-fidelity simulations guarantees unbiasedness of the estimation. To formalize the AMFMC implementation, let $\tilde{z}_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ denote the surge response approximation for the *i*th location obtained by the surrogate model and $\tilde{h}_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ denote the corresponding consequence measure. Let $N_{HF}^{(k)}$ and $N_{LF}^{(k)}$ represent the model evaluations for the high- and low-fidelity models for the current advisory, respectively. The MC sampling creates the sample set for the low-fidelity model, and the responses are estimated $\{\tilde{z}_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t)) : l = 1,...,N_{LF}^{(k)}\}$. For the high-fidelity model, the first $N_{HF}^{(k)}$ samples of this set are used to obtain its response predictions $\{z_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t)) : l = 1,...,N_{HF}^{(k)}\}$. The MFMC estimation is then formulated as[3]:

$$\hat{H}_{i,MF}^{(k)} = \frac{1}{N_{HF}^{(k)}}\sum_{l=1}^{N_{HF}^{(k)}} h_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t)) + a_i^{(k)}\left(\frac{1}{N_{LF}^{(k)}}\sum_{l=1}^{N_{LF}^{(k)}} \tilde{h}_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t)) - \frac{1}{N_{HF}^{(k)}}\sum_{l=1}^{N_{HF}^{(k)}} \tilde{h}_i(\mathbf{x}^l\,|\,\overline{\mathbf{q}}^{(k)}(t))\right) \quad (9)$$

where $a_i^{(k)}$ represents the control variate parameter for the *i*th QoI.

The two parameters that determine the statistical accuracy of the MFMC estimation in Eq. (9), are the control variate parameter, $a_i^{(k)}$ and the budget allocation between the high- and low-fidelity models, formally represented by the ratio $r^{(k)} = N_{LF}^{(k)} / N_{HF}^{(k)}$. Note that the selection of these parameters impacts only the variance of the MFMC estimator but not its bias. The optimal values for both $a_i^{(k)}$ and $r^{(k)}$ are the ones that minimize the variance of the estimator in Eq. (9), which can be performed in two stages. First, minimization with respect to $a_i^{(k)}$ leads to the following optimal control variate for the *i*th QoI[3]:

$$(a_i^{(k)})^* = \rho_i^{(k)}\sqrt{Var_p[h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))] / Var_p[\tilde{h}_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]} \quad (10)$$

where $\rho_i^{(k)}$ is the correlation between the high- and low-fidelity models for the QoI:

$$\rho_i^{(k)} = \frac{Cov_p[h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t)), \tilde{h}_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]}{\sqrt{Var_p[h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]\cdot Var_p[\tilde{h}_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))]}} \quad (11)$$

where $Cov_p[\cdot,\cdot]$ denotes the covariance operator under $p(\mathbf{x})$. With $(a_i^{(k)})^*$ in Eq. (10), the variance of the MFMC estimator becomes[3]:

$$Var[\hat{H}_{i,MF}^{(k)}] = \frac{Var_p[h_i(\mathbf{x}\,|\,\bar{\mathbf{q}}^{(k)}(t))]}{N_{HF}^{(k)}}\left(1-\left(1-\frac{1}{r^{(k)}}\right)\left(\rho_i^{(k)}\right)^2\right) \tag{12}$$

Further minimizing this variance with respect to $r^{(k)}$ leads to the following optimal ratio[3]:

$$(r_i^{(k)})^* = \sqrt{[c_{HF}\cdot(\rho_i^{(k)})^2]/[c_{LF}\cdot(1-(\rho_i^{(k)})^2)]} \tag{13}$$

where c_{HF} and c_{LF} represent the computational cost of the high- and low-fidelity models for a single simulation run, respectively. The optimal control variate parameter in Eq. (10) and the optimal budget allocation in Eq. (13) are estimated using a small number of pilot high-fidelity simulations and corresponding low-fidelity ones[3], before the MFMC estimation in Eq. (9).

Considering the challenge in the compromise between the QoIs, the optimal control variate parameter can be independently chosen for each QoI, since the choice for it only impacts the weight of the two different components of the estimate in Eq. (9). This is not true for the optimal budget allocation, since the same numerical simulations are used for the estimation of all QoIs. Therefore, a common budget allocation needs to be enforced. However, Eq. (13) indicates that each QoI promotes a different optimal budget allocation. Competing preferences among QoIs in the AMFMC framework are manifested as differences in the optimal budget allocations $(r_i^{(k)})^*$, originating from different correlation coefficients in Eq. (11). This is demonstrated in Figure 2, depicting the distribution over the domain of the correlation coefficient and the budget allocation for the same case as in Figure 1. Great variability is observed across the domain with different QoIs promoting different preferences for the budget allocation. To establish a well-balanced compromise solution for a budget allocation across all QoIs, $r^{(k)}$, a representative correlation coefficient between the two models was introduced in [11], which is obtained by the weighted average of the correlation coefficients across all QoIs in Eq. (11) as:

$$\rho^{(k)} = \sum_{i=1}^{n_z}\psi_i\rho_i^{(k)} \Big/ \sum_{i=1}^{n_z}\psi_i \tag{14}$$

where ψ_i represents the relative importance (i.e., weight) of the ith QoI. It was recommended in [11] to utilize the surge response variance for the weight, i.e., $\psi_i = Var[z_i(\mathbf{x}\,|\,\bar{\mathbf{q}}^{(k)}(t))]$, as this gives greater importance to locations with greater statistical variability. Finally, using the representative correlation coefficient in Eq. (13) leads to the representative optimal budget allocation[11], denoted as $(r^{(k)})^*$ herein. The detailed, computationally efficient AMFMC implementation for establishing $(r^{(k)})^*$ through dimensionality reduction can be found in [11].

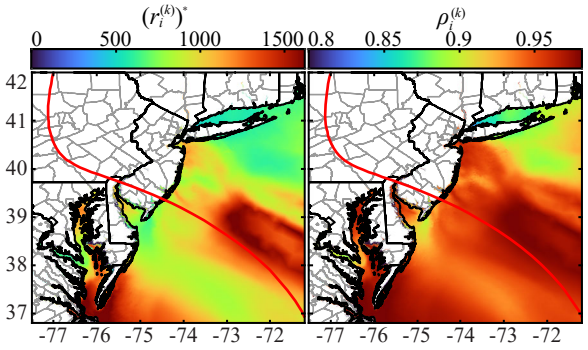

Figure 2: Spatial variability of the optimal budget allocation between the low- and high-fidelity models (left) and correlation coefficient between the two models (right) for the same case as in Figure 1.

It is important to note that the developed surrogate model and the selected representative budget allocation are uniformly applied across any surge statistics of interest in the MFMC estimation. The MFMC estimate of Eq. (9) is independently applied for each surge statistic of interest and each QoI (including estimation of optimal control variate for the QoI), but the computational implementation details are the same across all statistics.

3.3 Evaluation of the adaptive MC estimation efficiency

The comparison across the two adaptive MC frameworks is established by the speed-up measure, initially introduced in[13], to quantify efficiency improvement over the direct-MC, and by examining their idealized and practical implementations. The idealized implementation promotes the optimal adaptive choices independently for each QoI, while the practical one enforces a single, compromise choice across all QoIs discussed in the previous two subsections. The speed-up represents the computational time saved by each adaptive approach, compared to the direct-MC estimation, to achieve the same level of statistical accuracy.

The speed-up is estimated by the ratio of the variance of the direct-MC estimator in Eq. (3) to the variance of each of the adaptive estimators, using Eq. (5) for AIS or Eq. (12) for AMFMC, and will be denoted as $Sp_i^{(k)}$ for the ith QoI and the (k)th advisory. For AIS, it is expressed as:

$$Sp_{i,IS}^{(k)}[f^{(k)}(\mathbf{x})] = \frac{Var[\hat{H}_i^{(k)}]}{Var[\hat{H}_{i,IS}^{(k)}]} = \frac{E_p[h_i^2(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))] - (H_i^{(k)})^2}{E_{f^{(k)}}\left[\left(h_i(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))\frac{p(\mathbf{x})}{f^{(k)}(\mathbf{x})}\right)^2\right] - (H_i^{(k)})^2} \tag{15}$$

The *idealized* speed-up for each QoI is obtained by using the best density approximation for that QoI $f^{(k)}(\mathbf{x}) = \hat{f}_i^*(\mathbf{x}^I\,|\,\overline{\mathbf{q}}^{(k)}(t))$, with $\hat{f}_i^*(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ established base on simulation results from the current (k)th advisory. The *practical* speed-up is obtained by $f^{(k)}(\mathbf{x}) = \hat{f}_{DIS}(\mathbf{x}\,|\,\alpha_{DIS})$ with $\hat{f}_{DIS}(\mathbf{x}\,|\,\alpha_{DIS})$ established through the implementation presented in Section 3.1. For AMFMC, the speed-up is given by:

$$Sp_{i,MF}^{(k)}[r^{(k)}] = \frac{Var[\hat{H}_i^{(k)}]}{Var[\hat{H}_{i,MF}^{(k)}]} = \frac{c_{HF}}{c_{HF} + c_{LF}\cdot r^{(k)}}\left(1 - \left(1 - \frac{1}{r^{(k)}}\right)\cdot(\rho_i^{(k)})^2\right)^{-1} \tag{16}$$

The *idealized* speed-up for each QoI is obtained by using the best budget allocation for that QoI $r^{(k)} = (r_i^{(k)})^*$ given by Eq. (13). The *practical* speed-up is obtained by selecting as a compromise solution the representative ratio $r^{(k)} = (r^{(k)})^*$, obtained through the representative correlation coefficient presented in Section 3.2. Beyond the individual speed-up estimation for each QoI, the global efficiency improvement is examined through the weighted average of the speed-up across the domain of interest using the surge response variance as the weight.

4 ILLUSTRATIVE CASE STUDY

4.1 Case study characteristics and computational details

Case study considered implementation for two consecutive NHC advisories (25 and 26) of Superstorm Sandy (2012). ADCIRC is considered as the high-fidelity model for predicting storm surge, and surge statistics are obtained for a regional grid including n_z=1,860,021 nodes. To estimate reference results, especially to assess the *idealized* and *practical* speed-ups, a high-accuracy ADCIRC-based metamodel is utilized as high-fidelity model[4]. The estimation of

three different statistics is examined: the surge thresholds $b_i^{0.1}$ (corresponding to 10% exceedance probability) and $b_i^{0.01}$ (corresponding to 1% exceedance probability), and the mean surge. Notations $S_{0.1}$, $S_{0.01}$, and S_m will be used to distinguish them. The real-time probabilistic surge predictions are performed for the latter advisory (advisory 26). The previous advisory (advisory 25) is used to accommodate the information sharing. For AIS, this is the established formulation, with numerical results from a single previous advisory used to accelerate the MC estimation for the next one. For AMFMC, though, numerical results from multiple advisories can be utilized as discussed in Section 3.2. This is investigated by considering two different AMFMC implementations, utilizing results from one or two advisories for the development of the surrogate model. These will be distinguished by notations AMFMC$_s$ (single-advisory) and AMFMC$_m$ (multiple advisories), respectively. Note that the latter only utilizes results from two advisories to establish a more consistent comparison to AIS. In principle, more than two advisories could be used in the AMFMC formulation. In terms of computational details, for AIS, the IS densities are established using N=500 simulations, and separately selected for the three different surge statistics. For the practical implementation, the defensive scheme with γ=0.05 is utilized. For AMFMC, the number of pilot simulations is taken to be 20, following guidelines in [11], whereas a Gaussian process (GP)[14] is adopted as a surrogate model. The ratio of computational costs between the high- and low-fidelity models is $c_{HF}/c_{LF}=10^5$. The number of high-fidelity simulations available for each advisory for the development of the surrogate model is taken to be equal to the pilot estimations, meaning that surrogate models for AMFMC$_s$ and AMFMC$_m$ are established using 20 and 40 high-fidelity simulations, respectively.

4.2 Results and discussions

The global *idealized* and *practical* speed-ups are presented in Table 1. Focusing first on the idealized performance, the performance of AIS improves for lower likelihood events, while the performance of AMFMC deteriorates. This is expected, as AIS provides greater benefits for statistical products associated with rare events[10, 12]. In such cases, the difference between the $p(\mathbf{x})$ and $\hat{f}_i^*(\mathbf{x}\,|\,\overline{\mathbf{q}}^{(k)}(t))$ is bigger, enabling a larger variance reduction by IS implementation. On the other hand, the efficiency of AMFMC is expected to reduce for those statistical products[3], as the correlation between the high- and low-fidelity models typically reduces for samples away from the mean behavior[11]. Comparing the two AMFMC variants, it is evident that the ability to use numerical information from multiple advisories provides greater benefits: AMFMC$_m$ consistently outperforms AMFMC$_s$. From this point on, the AMFMC$_m$ variant will be taken as the representative AMFMC implementation and referenced simply as AMFMC.

Moving now to the *practical* speed-up, the comparisons show a significant performance degradation for AIS and a negligible one for AMFMC. The degradation for AIS is larger for the lower likelihood events, demonstrating a critical trend: the necessity to find a compromise solution across conflicting preferences for the QoIs, as demonstrated in Figure 1, poses greater challenges for the cases that AIS can accommodate higher efficiency (i.e., greater conflicts may arise). The need to infuse robustness to guarantee unbiasedness through the defensive IS scheme amplifies these challenges. On the other hand, since the correlation between the high- and low-fidelity models is generally high across the whole domain, as demonstrated in Figure 2, the representative correlation coefficient and budget allocation yield only a small performance deterioration for AMFMC. The greater versatility in sharing information across advisories and the smaller conflicts across QoIs promote ultimately a greater improvement for AMFMC.

Table 1: Weighted mean of the idealized and practical speed-ups for three statistics (S_m, $S_{0.1}$, and $S_{0.01}$) established by AIS and AMFMC (both variants) implementations.

speed-ups	AIS			AMFMC$_s$			AMFMC$_m$		
	S_m	$S_{0.1}$	$S_{0.01}$	S_m	$S_{0.1}$	$S_{0.01}$	S_m	$S_{0.1}$	$S_{0.01}$
idealized	10.22	6.03	46.22	24.63	5.05	2.43	29.13	7.83	2.54
practical	1.23	1.57	3.91	24.53	5.03	2.42	28.99	7.72	2.48

The spatial distribution of *idealized* (top row) and *practical* (bottom row) speed-ups is examined in Figures 3 and 4, respectively, for the AIS and AMFMC. Note that the scale is different for each subplot, chosen to better demonstrate the performance variability within the domain as well as the impact of the compromise solution on the spatial variability. Before continuing comparison discussions, it is important to note that the results in Figures 3 and 4 (and Table 1) for the *practical* speed-up, with quite high values, especially for AMFMC, showcase the actual benefits that can be achieved through the sharing of numerical information across advisories. This offers tangible proof to promote the use of such adaptive schemes.

Examining the trends associated with the difference between *idealized* and *practical* speed-ups across different statistical products in Figures 3 and 4, significant variability is observed. For AMFMC, results for the *idealized* and *practical* speed-ups are practically the same, showing that the compromise solution has a small impact, extending the trends already reported for the average speed-ups in Table 1. The speed-up distribution exhibits some mild similarities across the different statistical products, with subdomains with higher speed-up values for one statistical product having high values for other ones. This shows that the underlying mechanism for accelerating the MC estimation, the correlation between the high- and low-fidelity models, maintains consistent patterns across the QoIs for the different statistics.

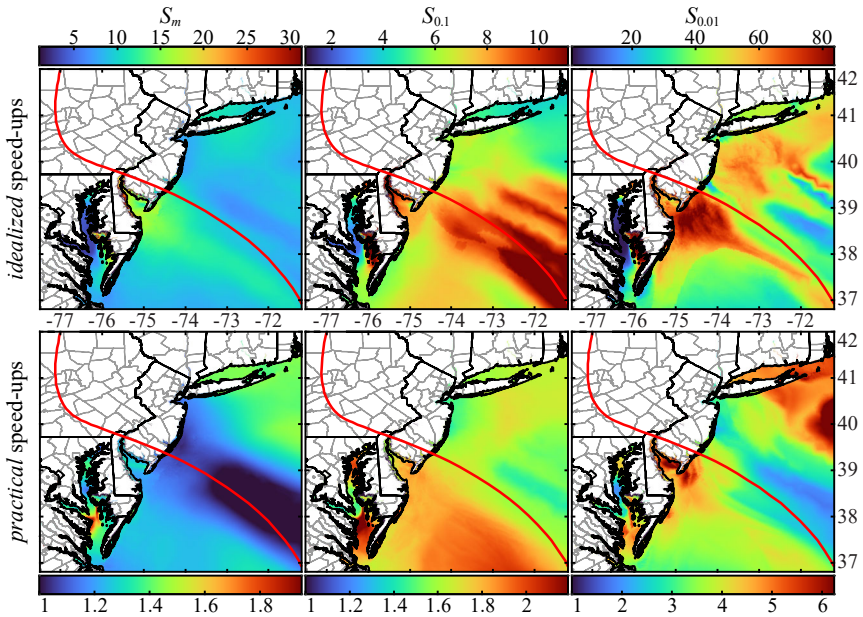

Figure 3: Spatial variability of idealized [top row] and practical [bottom row] speed-ups over the domain established by AIS with respect to statistics S_m [left column], $S_{0.1}$ [middle column], and $S_{0.01}$ [right column] for Superstorm Sandy (2012). The red line represents the storm track.

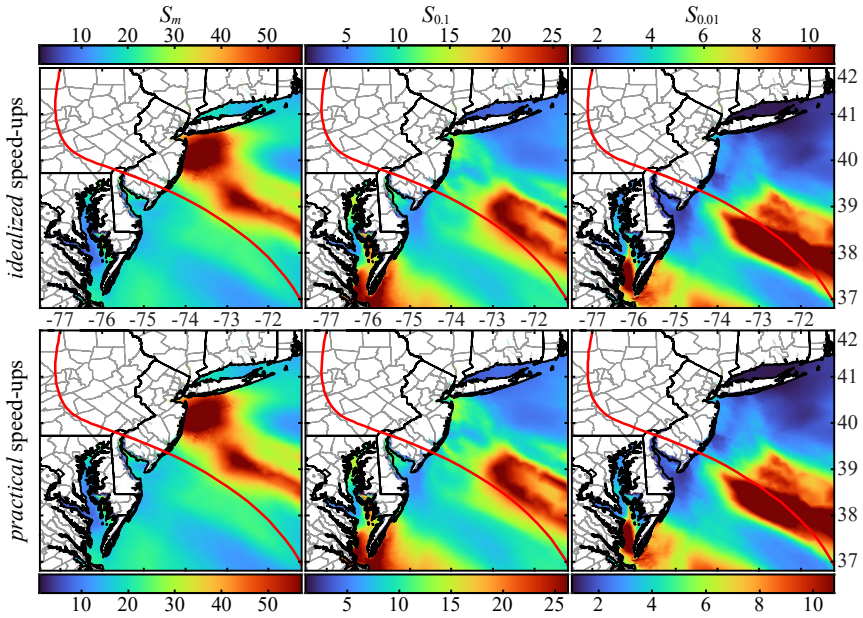

Figure 4: Spatial variability of idealized [top row] and practical [bottom row] speed-ups over the domain established by AMFMC with respect to statistics S_m [left column], $S_{0.1}$ [middle column], and $S_{0.01}$ [right column] for Superstorm Sandy (2012). The red line represents the storm track.

For AIS, the trends drastically differ from the ones identified for AMFMC. The variability of the speed-up distribution across the statistical products and across the idealized and practical formulations exhibits dramatic changes. The benefits from the AIS implementation are very different across the different surge statistics of interest. The enforcement of the compromise solution in the practical implementation fundamentally changes the distribution patterns within the domain. This is especially true for the $S_{0.01}$, as expected, since in this case, the AIS establishes greater performance improvements, and so is impacted at a greater degree by any conflicting preferences across the QoIs. These discussions further stress the challenges AIS can encounter in accommodating a compromise solution, as different parts of the domain compete with one another in the selection of IS densities for the different storm features.

5 CONCLUSIONS

This paper discussed the challenges in adopting AIS and AMFMC in regional risk assessment, examining a specific application, the real-time probabilistic storm surge predictions. It was shown that establishing a compromise solution of the adaptive MC characteristics for the high-dimensional QoIs (corresponding to storm surge statistics at different locations within the domain of interest) may pose a challenge, as QoIs promote different and competing preferences for them. The impact of these challenges on the performance of each adaptive technique was examined in detail through the comparison between its idealized and practical implementations. AIS exhibited a significant performance deterioration in its practical implementation compared to the idealized counterpart, whereas AMFMC showed only marginal differences between the two. These observations were demonstrated in both the averaged speed-ups across the geographic domain and their variability over the domain. The trends for the latter variability were altered for AIS across different surge

statistics and between the idealized and practical implementations, but were kept fairly similar for AMFMC. Overall the substantial efficiency improvements, particularly by AMFMC, also clearly demonstrate the computational benefits that can be achieved through the sharing of numerical information across advisories and promote the use of such adaptive schemes to support real-time probabilistic surge estimation to replace current practices of focusing on a single advisory for formulating the probabilistic predictions.

REFERENCES

[1] Deierlein GG, McKenna F, Zsarnóczay A, Kijewski-Correa T, Kareem A, Elhaddad W et al. A cloud-enabled application framework for simulating regional-scale impacts of natural hazards on the built environment. Frontiers in Built Environment. 2020;6:558706.

[2] Au SK, Beck JL. A new adaptive importance sampling scheme. Structural Safety. 1999;21:135-58.

[3] Peherstorfer B. Multifidelity Monte Carlo estimation with adaptive low-fidelity models. SIAM/ASA Journal on Uncertainty Quantification. 2019;7:579-603.

[4] Kyprioti AP, Taflanidis AA, Nadal-Caraballo NC, Campbell M. Storm hazard analysis over extended geospatial grids utilizing surrogate models. Coastal Engineering. 2021:103855.

[5] Hesterberg TC. Advances in importance sampling: Stanford University; 1988.

[6] Taylor AA, Glahn B. Probabilistic guidance for hurricane storm surge. 19th Conference on probability and statistics2008.

[7] Penny AB, Alaka L, Taylor AA, Booth W, DeMaria M, Fritz C et al. Operational storm surge forecasting at the National Hurricane Center: The Case for probabilistic guidance and the evaluation of improved storm size forecasts used to define the wind forcing. Weather and Forecasting. 2023;38:2461-79.

[8] Luettich RA, Jr. , Westerink JJ, Scheffner NW. ADCIRC: An advanced three-dimensional circulation model for shelves, coasts, and estuaries. Report 1. Theory and methodology of ADCIRC-2DDI and ADCIRC-3DL. Vicksburg,MS: Dredging Research Program Technical Report DRP-92-6, U.S Army Engineers Waterways Experiment Station; 1992.

[9] Kyprioti AP, Adeli E, Taflanidis AA, Westerink JJ, Tolman HL. Probabilistic Storm Surge Estimation for Landfalling Hurricanes: Advancements in Computational Efficiency Using Quasi-Monte Carlo Techniques. Journal of Marine Science and Engineering. 2021;9:1322.

[10] Jung W, Taflanidis AA, Kyprioti AP, Adeli E, Westerink JJ, Tolman H. Efficient probabilistic storm surge estimation through adaptive importance sampling across storm advisories. Coastal Engineering. 2023:104287.

[11] Jung W, Taflanidis AA, Kyprioti AP, Zhang J. Adaptive Multi-fidelity Monte Carlo for real-time probabilistic storm surge predictions. Reliability Engineering & System Safety. 2024:109994.

[12] Kroese DP, Taimre T, Botev ZI. Handbook of Monte Carlo methods. New York: John Wiley and Sons; 2011.

[13] Patsialis D, Taflanidis A. Multi-fidelity Monte Carlo for seismic risk assessment applications. Structural Safety. 2021;93:102129.

[14] Gramacy RB. Surrogates: Gaussian process modeling, design, and optimization for the applied sciences: Chapman and Hall/CRC; 2020.