Adaptive importance sampling for efficient probabilistic storm surge estimation

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

Aikaterini P. Kyprioti
Professor, Department of Civil Engineering and Environmental Science, University of Oklahoma, Norman, United States

ABSTRACT: During landfalling tropical storms, probabilistic predictions of the storm surge constitute important products for guiding emergency response decisions. The probabilistic formulation for these predictions is established by considering historical forecast errors for the intensity, size, cross- and along-track variability of the National Hurricane Center (NHC) advisories. These errors quantify ultimately uncertainties in storm features, serving as input to a numerical model for predicting storm surge, while propagation of these uncertainties provides the desired statistical products. This probabilistic estimation is repeated whenever the NHC updates the storm advisory. Monte Carlo (MC) simulation is considered for facilitating the uncertainty propagation in this paper, and in order to improve computational efficiency, the implementation of adaptive importance sampling (IS) across the storm advisories is introduced, using simulation results from the current advisory to select the optimal IS density to use for the next advisory. The requirement to estimate the storm surge across a large geographic domain, leading to the definition of a large number of quantities of interests (QoIs), poses a significant challenge, since these quantities typically represent competing IS choices. Principal component analysis (PCA) is utilized for dimensionality reduction to establish a compromising solution with a reduced computational burden. An adaptive selection of the IS characteristics is discussed, utilizing an efficient estimation of the anticipated IS accuracy, and a defensive IS scheme is introduced to guarantee robustness.

1. PROBLEM FORMULATION

A detailed description of the national weather service uncertainty quantification framework for probabilistic surge estimation can be found in (Kyprioti et al. 2021). Here a brief review is offered. During landfalling storms, the National Hurricane Center (NHC) provides the official advisory for the storm's: (i) forecasted track, described by the latitude, s_{lat}, and longitude, s_{lon} of its center; (ii) size, described by the radius of maximum winds, R_{mw}; (iii) and intensity, described by the maximum sustained wind speed, v_w. Let $\mathbf{q}(t)=[R_{mw}(t), s_{lat}(t), s_{lon}(t), v_w(t)]^T$ denote the four-dimensional vector of storm features, and t denote time. The advisory provides nominal predictions of those features, denoted herein as $\tilde{\mathbf{q}}(t)$. The NHC also provides forecast errors associated with each advisory for four different characteristics: the cross- and along-track, Δs_{cross} and Δs_{along}, respectively, the storm size, ΔR_{mw}, and the storm intensity, Δv_w. The vector characterizing these variabilities is defined as $\boldsymbol{\Delta}\mathbf{q}(t)=[\Delta R_{mw}(t), \Delta s_{cross}(t), \Delta s_{along}(t), \Delta v_w(t)]^T$.

The statistical analysis of past forecast errors provides a probabilistic description of $\boldsymbol{\Delta}\mathbf{q}(t)$. The current NHC formulation assumes independence between the four errors and perfect correlation across times (Taylor and Glahn 2008). This simplifies the probabilistic description, requiring simply a four-dimensional vector of random variables, denoted herein by $\mathbf{x} \in \mathbb{R}^{n_x}$ (n_x=4), for the uncertainty description. Details for the probability models for **x**, denoted as $p(\mathbf{x})$, are included in (Kyprioti et al. 2021). In the implementation discussed herein transformation to standard Gaussian space is adopted for $p(\mathbf{x})$.

Vector **x**, defines ultimately the variability vector $\Delta\mathbf{q}(t)$, which can be then combined with the nominal advisory information to provide a sample storm scenario $\mathbf{q}(t)$. This scenario can be then used as input to an appropriate hydrodynamic model to predict the storm surge across n_z discretized locations within the geographic domain the storm impacts. Let $z_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})$ (i=1,…,n_z) denote the surge at a location i and use $h[z_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})] = h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})$ to express the consequence measure related to the statistic of interest, with the latter expressed as:

$$H_i = E_p[h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})] = \int h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})p(\mathbf{x})d\mathbf{x} \quad (1)$$

where E_p[.] is the expectation under $p(\mathbf{x})$. For example for the probability $P_i(b)$ that the surge will exceed a specific threshold b, $h[z_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})]$ is equal to $I[z_i(\mathbf{x}\,|\,\tilde{\mathbf{q}}) > b]$, where I[.] is the indicator function. The common statistical product used by NHC to convey hazard is the surge value $b_i^{p_t}$ corresponding to a probability of exceedance p_t, so that $P_i(b_i^{p_t}) = p_t$.

2. MC AND IMPORTANCE SAMPLING

This contribution considers the estimation of the statistics of interest though Monte Carlo estimation (MC), leveraging importance sampling (IS) for efficiency. Considering i.i.d samples {$\mathbf{x}^l$: l=1,…,N} from proposal density $f(\mathbf{x})$, the MC estimator, $\hat{H}_i$ for the integral in Eq. (1) is:

$$\hat{H}_i(f) = \frac{1}{N}\sum_{l=1}^{N} h_i(\mathbf{x}^l\,|\,\tilde{\mathbf{q}})p(\mathbf{x}^l)/f(\mathbf{x}^l) \quad (2)$$

Provided that the integrand $h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})p(\mathbf{x})/f(\mathbf{x})$ is bounded, estimator $\hat{H}_i$ is unbiased with coefficient of variation, δ_i, representing its statistical accuracy as an approximation of H_i, given by (Kroese et al. 2011):

$$\delta_i(f) = \frac{1}{\sqrt{N}} \frac{\sqrt{Var_f[h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})p(\mathbf{x})/f(\mathbf{x})]}}{H_i} \quad (3)$$

where Var_f[.] denotes variance under $f(\mathbf{x})$.

Direct MC uses $f(\mathbf{x})=p(\mathbf{x})$, while for MC with IS $f(\mathbf{x})$ is chosen so that the sampling concentrates in regions corresponding to larger values for the integrand in Eq. (1). The IS objective is to reduce the variance appearing in the numerator of Eq. (3) while the optimal IS density, establishing the minimization, is (Kroese et al. 2011):

$$f_i^*(\mathbf{x}\,|\,\tilde{\mathbf{q}}) \propto |h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}})|\,p(\mathbf{x}) \quad (4)$$

Using directly this density is impractical since it requires knowledge of the entire integrand. A sample-based approximation of this density will be considered here instead, incorporating adaptive selection of some of its features. This is achieved by examining different candidates $f^c(\mathbf{x})$ and promoting the one with the smallest anticipated efficiency (Medina and Taflanidis 2014).

3. ADAPTIVE IS FOR PROBABILISTIC STORM SURGE ESTIMATION

To implement the sample-based IS formulation for the probabilistic surge estimation, three challenges need to be considered: (i) the necessity to choose the sample-based proposal densities using limited information since the number of simulations N is typically small due to the large computational burden of the numerical models; (ii) the need to establish a compromise across conflicting outputs with each of them leading to a different optimal density; (iii) the requirement to share information across the storm advisories, basing the selection of the proposal density for the next advisory on the responses available for the current advisory.

The adaptive IS framework is established by choosing IS densities utilizing the storm surge simulations from the current (k) advisory, with the intension to use these simulations for the next (k+1) advisory. A superscript in parenthesis $.^{(k)}$ will be used hereinafter to explicitly denote the kth storm advisory. Responses for the current advisory $\{h_i(\mathbf{x}^l\,|\,\tilde{\mathbf{q}}^{(k)}) : l = 1,...,N\}$ obtained using sample set $\{\mathbf{x}^l : l = 1,...,N\} \sim f^{(k)}(\mathbf{x})$ are readily available to support the IS selection. We will denote as $\varphi(\mathbf{x}) \sim p(\mathbf{x}) / f^{(k)}(\mathbf{x})$ the IS-ratio for the current advisory. Let also $\mathbf{H}^{(k)} \in \mathbb{R}^{N \times n_z}$ denote the observation matrix with rows corresponding to the QoI consequence measure values for the lth storm simulation.

3.1. Sample-based approximation of IS density

Let $\pi(\mathbf{x})$ denote the target optimal density (this will be defined in next section). Based on resampling principles, the currently available samples $\{\mathbf{x}^l : l = 1,...,N\} \sim f^{(k)}(\mathbf{x})$ provide information for the updated target density $\pi(\mathbf{x})$ if weight $w(\mathbf{x}) = \pi(\mathbf{x}) / f^{(k)}(\mathbf{x})$ is attached to each of the samples. Utilizing the pair of samples/weights $\{(\mathbf{x}^l, w^l): l=1,\ldots,N\}$, a sample-based approximation for $\pi(\mathbf{x})$ can be established by fitting some chosen distribution. A parametric density approximation is adopted here to accommodate the need to establish this approximation using limited available information. i.e. small number of simulations \. Specifically, a Gaussian Mixture Model (GMM) approximation is selected. The GMM fit is established through Expectation-Maximization with the weights $w(\mathbf{x})$ incorporated in the formulation through an adjustment of the data-likelihood as suggested in (Geyer et al. 2019).

To improve the robustness of the distribution fit, both marginal and joint density formulations are examined, and, additionally, the inclusion of the most influential input components in the IS formulation is considered. The motivation for examining these alternative options is to further address overfitting issues associated with the limited available information. The prioritization of input components or the use of marginal densities promotes extraction of smaller amount of information from the sample set, which can facilitate greater robustness in the IS selection: different candidate IS formulations are established, then the efficiency for each of them is approximated and the most appropriate one is finally chosen. For the input prioritization, a recently proposed GSA (Jung et al. 2022) is leveraged. This GSA provides aggregated importance indicators for input $\mathbf{x}$, quantifying the importance of each component towards the observed surge variability. The IS formulation can focus on only the important components. If $\mathbf{x}_s$ denotes the important components and $\mathbf{x}_{\sim s}$ the remaining ones, then the candidate would be $f^c(\mathbf{x}) = f^c(\mathbf{x}_s) p(\mathbf{x}_{\sim s})$. The sample-based approximation of $f^c(\mathbf{x}_s)$ can be established for the entire vector $\mathbf{x}_s$ (joint formulation) or for each of its components (marginal formulation). The terms IS$_M$ (marginal) and IS$_T$ (joint) will be used herein to distinguish these two implementations. Relevant mathematical and computational details can be found in (Jung et al. 2023).

3.2. Selection of IS densities

Each of the examined QoIs has a different optimal density, $f_i^*(\mathbf{x} | \tilde{\mathbf{q}}^{(k)})$, while, more importantly, these densities can be conflicting. Figure 1 illustrates this challenge, showing the mean for $f_i^*(\Delta s_{cross} | \tilde{\mathbf{q}}^{(k)})$ for superstorm Sandy, for two different advisories and estimation of 10% exceedance probability statistical products. This mean represents an important property for the IS density, as it depicts the shift towards the peak of the probabilistic integrand. It is evident that great variability of the mean exists within the geographic domain of storm impact, showcasing that the different QoIs (surge at different geographic locations) advocate conflicting decisions. The conflicts are expected to reduce the overall IS efficiency, since a single compromising IS density can be quite different from the actual optimal density for each QoI. Moreover, the selection of this density that balances the conflicting objectives needs to be carefully executed, while also accommodating applications with a large dimensional output.

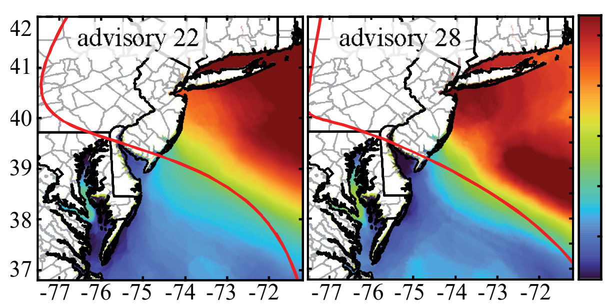

Figure 1: Variation of optimal IS density mean for Δs_{cross} for two advisories (22 and 28) of superstorm Sandy for estimation of 10% probability statistics.

To accommodate these requirements, an IS formulation that utilizes Principal Component Analysis (PCA) as a dimensionality reduction technique is considered here. Within this context, PCA identifies a low-dimensional principal component vector, that best explains the variability of the original data $\mathbf{H}^{(k)}$. To incorporate the weight that each observation is associated with, $\{\varphi(\mathbf{x}^l)$: $l=1,\ldots,N\}$, a weighted (generalized) PCA is implemented (Jolliffe 2002). Each principal component $\{y_v$: $v=1,\ldots,n_v\}$ has an associated importance λ_v, representing the variance of the data explained by the component. To accommodate the dimensionality reduction, only the more important components are retained, with the exact number chosen so that significant portion of the original output variance can be explained (over 99%). PCA provides also the responses across the simulations $\{y_v(\mathbf{x}^l)$: $l=1,\ldots,N\}$ for each of the principal components.

The optimal IS density for each component is established as:

$$f_v^*(\mathbf{x}^l \mid \mathbf{H}^{(k)}) \propto | y_v(\mathbf{x}^l) | \, p(\mathbf{x}^l) / \sqrt{\varphi(\mathbf{x}^l)} \qquad (5)$$

where the division by $\sqrt{\varphi(\mathbf{x}^l)}$ is incorporated to address the weights introduced in the generalized PCA (Jung et al. 2023). Each density given by Eq. (5) represents a different target density $\pi_v(\mathbf{x})$ that can be approximated through the sample-based approach discussed in Section 3.1. Since the number of such densities is small (value of n_v is small), a separate fitting to each of these densities is examined. This is the advantage offered by PCA, allowing the sample-based approximation to focus on each of the (small number) principal components instead of the original large dimensional QoI vector. Alternative formulations can be found in (Jung et al. 2023).

The final candidate IS density is obtained by a weighted sum of the individualized IS densities for each principal component (Hesterberg 1995), with weights corresponding to λ_v. This leads to the following candidate density definition:

$$f^c(\mathbf{x} \mid \mathbf{H}^{(k)}) = \frac{\sum_{v=1}^{n_v} \lambda_v \hat{f}_v^*(\mathbf{x} \mid \mathbf{H}^{(k)})}{\sum_{v=1}^{n_v} \lambda_v} \qquad (6)$$

The efficiency of the IS proposal density for each of the different outputs is characterized by the variance appearing in Eq. (3). Across the different outputs the average of this quantity can be utilized as an efficiency measure. Using as normalization the average MC estimator variability (without IS), the following efficiency measure is defined for the kth advisory:

$$F^{(k)}[f^c(\mathbf{x})] = 1 - C$$
$$C = \frac{\sum_{i=1}^{n_z} Var_{f^c}[h_i(\mathbf{x} \mid \tilde{\mathbf{q}}^{(k)}) p(\mathbf{x}) / f^c(\mathbf{x})]}{\sum_{i=1}^{n_z} Var_p[h_i(\mathbf{x} \mid \tilde{\mathbf{q}}^{(k)})]} \qquad (7)$$

This efficiency represents the reduction of computational effort attained by the proposal density $f^c(\mathbf{x})$ for establishing the same average statistical accuracy (same average variance) across the outputs as the direct MC estimation.

To accommodate an adaptive IS implementation, comparing across different candidate choices for $f^c(\mathbf{x})$, for example the IS$_M$ (marginal) and IS$_T$ (joint) formulations discussed in previous section of the PCA-based implementation, or the alternative formulations discussed (Jung et al. 2023), an approximation for the anticipated efficiency needs to be obtained using readily available information. Following implementation in (Medina and Taflanidis 2014), and using the readily available samples $\{\mathbf{x}^l : l = 1,\ldots,N\} \sim f^{(k)}(\mathbf{x})$, an MC-based estimate, termed *efficiency approximation*, is given by :

$$\hat{F}^{(k)}[f^c(\mathbf{x})] = 1 - \frac{1}{n_z}\sum_{i=1}^{n_z} A_i \Bigg/ \frac{1}{n_z}\sum_{i=1}^{n_z} B_i$$
$$A_i = \frac{1}{N}\sum_{l=1}^{N}\left(h_i(\mathbf{x}^l \mid \tilde{\mathbf{q}}^{(k)})\right)^2 \frac{p(\mathbf{x}^l)}{f^c(\mathbf{x}^l)}\frac{p(\mathbf{x}^l)}{f^{(k)}(\mathbf{x}^l)} - \left(\frac{1}{N}\sum_{l=1}^{N} h_i(\mathbf{x}^l \mid \tilde{\mathbf{q}}^{(k)}) \frac{p(\mathbf{x}^l)}{f^{(k)}(\mathbf{x}^l)}\right)^2 \qquad (8)$$
$$B_i = \frac{1}{N}\sum_{l=1}^{N}\left(h_i(\mathbf{x}^l \mid \tilde{\mathbf{q}}^{(k)})\right)^2 \frac{p(\mathbf{x}^l)}{f^{(k)}(\mathbf{x}^l)} - \left(\frac{1}{N}\sum_{l=1}^{N} h_i(\mathbf{x}^l \mid \tilde{\mathbf{q}}^{(k)}) \frac{p(\mathbf{x}^l)}{f^{(k)}(\mathbf{x}^l)}\right)^2$$

3.3. Robustness of IS densities across advisories

In the proposed implementation the selection of proposal densities is based on the current, kth advisory, but is employed in the next (k+1)th advisory. Loss of efficiency, but most importantly loss of robustness, may arise if the promoted IS density does not sufficiently cover the important regions (for efficiency) or support (for robustness) of the updated integrand $h_i(\mathbf{x}\,|\,\tilde{\mathbf{q}}^{(k+1)})p(\mathbf{x})$. Robustness is a bigger concern here, since it might lead to biased estimates. A defensive IS formulation is adopted to improve robustness. If $f_{IS}^{c}(\mathbf{x})$ denotes the IS density approximation, then the defensive IS, $f_{DIS}^{c}(\mathbf{x}\,|\,a)$, is obtained as (Hesterberg 1995):

$$f_{DIS}^{c}(\mathbf{x}\,|\,\alpha) = \alpha f_{IS}^{c}(\mathbf{x}) + (1-\alpha)p(\mathbf{x}) \qquad (9)$$

with $0 \leq \alpha \leq 1$ representing the defensive IS weight. The choice of this weight needs to be established with care, since small values of α can lead to a significant reduction in IS efficiency. An adaptive selection is suggested here for α by comparing the (reduced) efficiency $\hat{F}^{(k)}(f_{DIS}^{c})$ of $f_{DIS}^{c}(\mathbf{x})$ to the (optimal) efficiency $\hat{F}^{(k)}(f_{IS}^{c})$ of $f_{IS}^{c}(\mathbf{x})$, and selecting an appropriate defensive weight α^c so that the resulting reduction is less than a threshold ρ. This choice ensures the greater robustness and the corresponding density $f_{DIS}^{c}(\mathbf{x}\,|\,\alpha^{c})$ will be denoted as *robust-IS*. Note that as discussed in (Jung et al. 2023) this defensive IS scheme enhances robustness not only for the sharing information across the storm advisories, but additionally for accommodating decisions using small sample sets across a large number of QoIs. In some cases, the defensive IS might outperform the original IS. For this reason, an improvement of the optimal IS is first established, selecting α to maximize the *efficiency approximation*:

$$\alpha^{*} = \arg\max_{\alpha\in[0,1]} \hat{F}^{(k)}[f_{DIS}^{c}(\mathbf{x}\,|\,\alpha)] \qquad (10)$$

The $f_{DIS}^{c}(\mathbf{x}\,|\,\alpha^{*})$ is ultimately the promoted optimal IS (will be denoted as *promoted-IS*), and then a further decrease in α is considered to provide the *robust-IS* .

3.4. Adaptive implementation across advisories

Combining the earlier concepts, the adaptive IS formulation across the advisories, termed AIS-SA, is established as follows. (A) based on the readily available information from the current advisory, different candidate IS proposal densities are constructed in Eq. (6) utilizing the sample-based optimal IS densities of Eq. (5) for individual principal components. The candidates pertain to joint and marginal formulations, for different number of important variables within $\mathbf{x}$ (definitions of $\mathbf{x}_s$). (B) the anticipated efficiency of each candidate density is assessed by calculating the *approximated efficiency* of Eq. (8) whereas the potential efficiency improvement is incorporated through Eq. (10) to obtain the *promoted-IS* for each candidate choice. (C) the efficiency of the different candidate densities is compared and the one with the best efficiency is selected. (D) a defensive IS scheme is further adopted through Eq. (9) allowing a chosen reduction of efficiency ρ for accomodating the desired robustness improvement. (E) This obtained *robust-IS* density is incorporated in the MC implementation in the next advisory. Further details are included in (Jung et al. 2023).

4. CASE STUDY

The proposed IS formulation is showcased for superstorm Sandy (2012), considering estimation across NHC advisories 22-28, corresponding to a period range roughly from 72hr to 24hr before storm makes landfall. The earliest advisory (with the smallest indexing number), corresponding to the furthest time from landfall, will be denoted by $A^{(1)}$ and each subsequent advisory will be denoted with an increasing superscript number. The domain of storm impact includes n_z=1,860,021 nodes, creating a high-dimensional output. Unless otherwise specified, the number of samples for the MC and IS implementations is chosen as N=500, whereas the efficiency reduction for the *robust-IS* selection is set to ρ=0.1. Implementation for two statistical products is examined, corresponding to estimation of the probability of exceeding thresholds $b_i^{0.1}$ and $b_i^{0.01}$. These two cases will be denoted as $S_{0.1}$ and $S_{0.01}$, respectively. For validation purposes, reference results for the statistical accuracy are obtained using quasi-Monte Carlo (QMC) with a large number (N=5000) of samples.

4.1. Adaptive IS selection for a single advisory

Discussions first focus on the potential IS benefits, investigating a single advisory. The benefits across advisories will be examined in the next section. Results are presented for advisory $A^{(1)}$ for IS densities considering different number of input components for both the marginal and joint IS formulations. For notational simplicity, each of the densities will be denoted herein as M_n [for the marginal IS_M formulation] or T_n [for the joint IS_T formulation], where subscript $._n$ denotes the number of input components considered in the IS formulation (number of components for $\mathbf{x}_s$ definition). Initially, the discussion focuses on the *promoted-IS* density.

The efficiency results for both examined statistics are presented in Figure 2 for different formulations of the IS density (discrete cases in the x-axis). Both the anticipated efficiency (termed *approximation*) established through Eq. (8) as well as the reference results obtained through the QMC implementation (termed *actual*) are reported in this figure, utilizing different marker styles and colors. To examine additionally the impact of the amount of available information on the adaptive IS selection, Figure 3 replicates Figure 2 for the $S_{0.1}$ statistic using a smaller number of N=100 simulations.

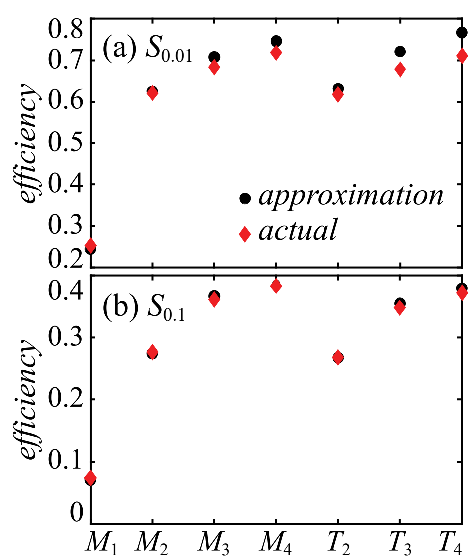

Figure 2: Approximated and actual IS efficiency for different IS formulations corresponding to $S_{0.01}$ (Top) and $S_{0.1}$ (Bottom) statistics.

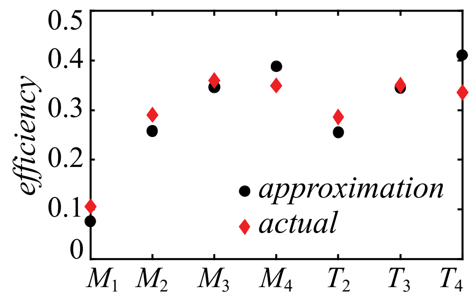

Figure 3: Approximated and actual IS efficiency for different IS formulations corresponding to $S_{0.1}$ statistic using N=100 simulations.

Focusing on the overall IS efficiency first, we can observe significant accuracy improvements, and proportional reduction in computational burden (1.5 to 2 times). The benefits are greater, as expected, for statistics corresponding to less frequent events; this is showcased here by a bigger improvement for $S_{0.01}$ compared to $S_{0.1}$. Examining the differences in the efficiency as the number of IS input variables increases (for both the marginal M_n or joint formulations T_n), n=4 emerges as the preferable implementation. To better frame the influence of the number of IS input variables, the GSA results need to be considered. The GSA identified Δs_{cross} as the most influential input (Jung et al. 2022). The results show the larger relative efficiency improvement when moving from n=1 to n=2, as opposed to when moving from n=0 (no IS implementation) to n=1 (including only Δs_{cross} in the IS formulation). Despite Δs_{cross} being the dominant input with respect to global sensitivity characteristics, it is not the most advantageous to include in the IS formulation. Looking at Figure 1 again, optimal IS characteristics for Δs_{cross} exhibit high variability within the domain, demonstrating clear competing preferences. The competing optimal IS characteristics and the need to establish a compromise across them, reduce its relative influence in the IS formulation.

Examining the behavior of the marginal IS_M and joint IS_T density formulations across Figures 2 and 3 one immediately observes similar performance, with the marginal implementation IS_M exhibiting some slightly better performance. The better performance for the IS_M family of densities is more evident when examining the *actual efficiency* for the implementation that has a smaller number of simulations (Figure 3), something that translates to a limited amount of information for selecting the adaptive IS. This is manifested as some (slight) loss of robustness for the joint family of densities, especially when the number of input variables n considered in the IS implementation is larger.

Comparing, finally, the approximated and actual efficiency, good overall agreement is reported. The discrepancies are significantly bigger, as expected, for (i) the case corresponding to less frequent event [comparing between cases in Figure 2], since a smaller amount of information is extracted from the same amount of resources (same number of available simulations), and (ii) for the case with smaller N [comparing between Figures 2 and 3], owing to the limited information utilized in this case to obtain the *efficiency approximation*. Nevertheless, the adaptive IS selection based on the *efficiency approximation* yields consistent results with the choices that would be made if the *actual efficiency* was known. This is the most important feature, as the *efficiency approximation* is directly leveraged within the AIS-SA to select the best IS candidates.

4.2. IS formulation across advisories

Comparisons move next to the implementation across advisories. To evaluate the efficiency of the information sharing across advisories, another formulation is examined, corresponding to the optimal IS selection for the current advisory. This allows the efficiency accomplished by choosing IS densities across subsequent advisories to be compared to the best efficiency that would be achieved if one could explicitly choose and implement the IS density for the present advisory. Terminology *previous* will be used to represent the AIS-SA formal implementation (across advisories) and *current* to represent the best efficiency. Results for the actual efficiency, obtained through the high-accuracy QMC implementation, are reported for both the *promoted-IS* and *robust-IS* densities. Figure 4 presents results for the $S_{0.01}$ and the $S_{0.1}$ statistics for the original implementation of N=500, while Figure 5 examines the implementation for a smaller number of N=100 for the $S_{0.1}$ statistic.

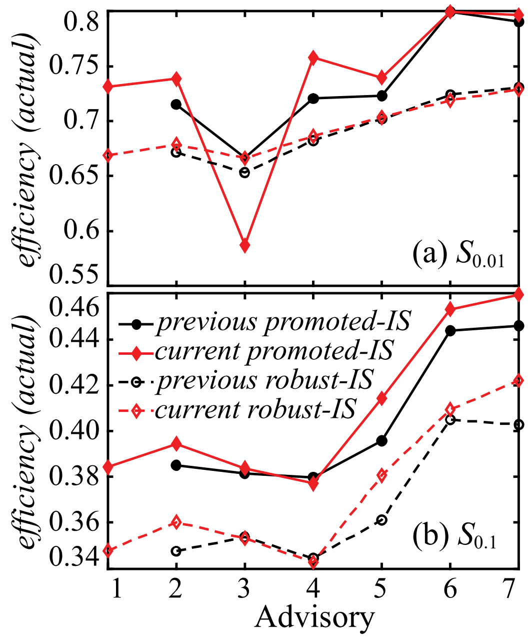

Figure 4: Efficiency of AIS-SA across advisories for the $S_{0.01}$ (Top) and $S_{0.1}$ (Bottom) statistics for both the promoted-IS density and the robust-IS density.

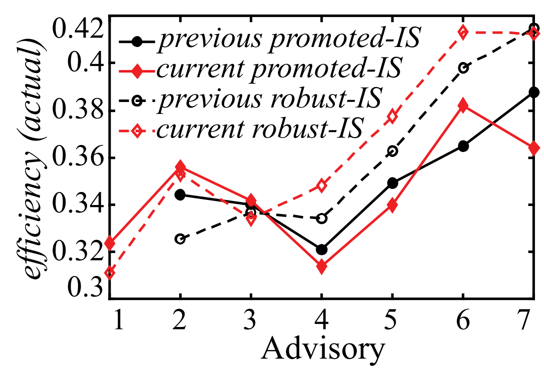

Figure 5: Efficiency of AIS-SA across advisories for the $S_{0.1}$ statistic for both the promoted-IS density and the robust-IS density using N=100 simulations.

Results clearly show that the proposed implementation, sharing information across advisories, works well, since the gap between the *previous* and *current* IS implementations is small. Note that the fact that in some (very few) instances the *previous IS* outperforms the *current IS*, is not something that should be anticipated but it can be explained: the selection of the IS densities is based on the *efficiency approximation* but the validation here is performed with respect to the *actual efficiency*. There can be instances that the *efficiency approximation* leads to (slightly) sub-optimal decisions when evaluated against the *actual efficiency* since it is based on limited information (number of simulations).

Investigating further trends in Figures 4 and 5, the implementation of the defensive IS strategy (*robust-IS* case) does not seem to be warranted from the perspective of sharing information across advisories, since the formulation based on the previous advisory information is very close to the *optimal* IS selection, indicating no loss of robustness by the fact that advisories change. It is warranted though to provide robustness against making selections using limited information for deciding on the IS densities (N=100), as the results in Figure 5 clearly indicate. As discussed earlier, the *robust-IS* addresses potential vulnerabilities that the AIS-SA algorithm could have when it selects IS density based on the previous advisory with limited information by evaluating IS candidates using their approximated efficiency. Since maintaining a proper level of robustness in the AIS-SA formulation is important, depending on the number of samples available, ρ could be adjusted appropriately. If N is large, smaller ρ can be utilized based on the results in Figures 4-5. Unfortunately, an adaptive way to choose ρ cannot be established since the detailed exploration of an appropriate value of ρ requires estimation of the *actual efficiency* for the next advisory.

5. CONCLUSIONS

This paper examined the importance sampling (IS) implementation for improving computational efficiency in real-time probabilistic storm surge estimation. The IS is established by sharing information across estimation for different advisories, using the readily available simulation results from the current advisory to adaptively select a proposal density, to use for the next advisory. A sample-based, adaptive IS formulation, AIS-SA, was established to achieve the objective. The case study showed that information sharing can facilitate satisfactory computational benefits, though challenges exist associated with the conflicting IS priorities set by the different QoIs. Additionally, the differences between subsequent advisories are, typically, not significant enough to lead to erroneous decisions within the proposed information sharing. However, the additional robustness established through the defensive-IS scheme is critical when the available information for formulating the IS densities across advisories is limited.

6. ACKNOWLEDGMENTS

This research was funded by the National Oceanographic and Atmospheric Administration (NOAA), under the grant number NA19OAR0220089. The views and opinions expressed in this paper are those of the authors and do not represent NOAA's official position.

7. REFERENCES

Geyer, S., Papaioannou, I., and Straub, D. (2019). "Cross entropy-based importance sampling using Gaussian densities revisited", *Structural Safety*, 76, 15-27.

Hesterberg, T. (1995). "Weighted average importance sampling and defensive mixture distributions", *Technometrics*, 37(2), 185-194.

Jolliffe, I.T. (2002). *Principal component analysis*, Springer series in statistics, 2nd edn. Springer, New York.

Jung, W., Kyprioti, A.P., Adeli, E., and Taflanidis, A.A. (2022). "Exploring the sensitivity of probabilistic surge estimates to forecast errors", *Natural Hazards*, 1-39.

Jung, W., Taflanidis, A.A., Kyprioti, A.P., Adeli, E., Westerink, J.J., and Tolman, H. (2023). "Efficient probabilistic storm surge estimation through adaptive importance sampling across storm advisories", *Coastal engineering*, in press.

Kroese, D.P., Taimre, T., and Botev, Z.I. (2011). *Handbook of Monte Carlo methods*, Wiley Series in Probability and Statistics, John Wiley and Sons, New York.

Kyprioti, A.P., Adeli, E., Taflanidis, A.A., Westerink, J.J., and Tolman, H.L. (2021). "Probabilistic Storm Surge Estimation for Landfalling Hurricanes: Advancements in Computational Efficiency Using Quasi-Monte Carlo Techniques", *Journal of Marine Science and Engineering*, 9(12), 1322.

Medina, J.C., Taflanidis, A.A. (2014). "Adaptive importance sampling for optimization under uncertainty problems", *Computer Methods in Applied Mechanics and Engineering*, 279, 133–162.

Taylor, A.A. and Glahn, B. (2008). "Probabilistic guidance for hurricane storm surge", *In 19th Conference on probability and statistics*.