https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

UPSurgeML (v1.0) - A Machine Learning Workflow for Probabilistic Ensemble Forecast of Tropical Cyclone Storm Surge

Fariborz Daneshvar[1,2], Soroosh Mani[1], William J. Pringle[3], Panagiotis Velissariou[4], Zizang Yang[1], Gregory Seroka[1], Edward Myers[1], Saeed Moghimi[1]

[1] Coast Survey Development Laboratory, Office of Coast Survey, NOAA National Ocean Service, Silver Spring, MD, USA

[2] Ocean Associates Inc. (OAI), Arlington,VA, USA

[3] Environmental Science Division, Argonne National Laboratory

[4] Cooperative Programs for the Advancement of Earth System Science (CPAESS), University Corporation for Atmospheric Research (UCAR), Boulder, CO, USA

Correspondence: Fariborz Daneshvar (fariborz.daneshvar@noaa.gov) and Soroosh Mani (soroosh.mani@noaa.gov)

Abstract

UPSurgeML is an open-source software package for probabilistic storm surge analysis. The code is written in Python, C, and Fortran, and the codebase complies with the National Centers for Environmental Prediction's (NCEP) guidelines, making it suitable for deployment on High Performance Computers (HPC) like the Weather and Climate Operational Supercomputing System (WCOSS). By utilizing an unstructured mesh, UPSurgeML refines the prediction of water elevations in complex coastal regions. It provides both deterministic and probabilistic storm surge forecasts driven by a small ensemble of hurricane forecasts. UPSurgeML uses parametric hurricane wind and pressure fields, and astronomical tide forcing with a hydrodynamic ocean model to simulate water elevations. Unlike the current Probabilistic Surge (PSurge) model utilized by the National Hurricane Center (NHC), UPSurgeML employs a machine learning surrogate

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

modeling framework. This approach enables a smaller (order of magnitude reduction) hydrodynamic ensemble size that would otherwise be necessary. Consequently, UPSurgeML provides high-resolution water elevation probability fields and exceedance levels with reasonably low computational cost, and has demonstrated its usability as experimental guidance during the hurricane seasons.

1. Introduction

Tropical cyclones (TCs) pose a significant threat to coastal communities (with more than 2,500 fatalities since 1963 in the U.S. Atlantic and Gulf coasts region), underscoring the critical need for accurate TC-driven coastal hazard forecasting. The National Oceanic and Atmospheric Administration (NOAA) National Hurricane Center (NHC) has been the U.S. pioneer of TC-driven coastal hazard assessment since the 1980s. NHC developed the Sea, Lake, and Overland Surge for Hurricanes (SLOSH) model in the early 1980s to simulate storm surge (Jelesnianski et al., 1992). SLOSH is a computationally efficient model that solves the linear shallow water equations on structured curvilinear grids. In the early 2000s, NHC developed the Probabilistic Surge (PSurge) model in order to take into account storm trajectory, size and intensity uncertainties in surge prediction (Glahn et al., 2009). PSurge uses historical forecast errors to generate a large ensemble (~400-1200 members) of perturbed storm tracks, and then uses SLOSH to simulate their corresponding surges for further probabilistic prediction of storm surge (PSurge, 2026). Recent advancements in prediction of the radius of maximum wind (RMW) and inclusion of spatially varying bottom friction overland has also improved PSurge forecast reliability (Penny et al., 2023; PSurge, 2026). However, some limitations regarding the small domain size and linearized form of the equations used by SLOSH have been highlighted in previous studies (Zhang et al., 2008; Kerr et al., 2013; Pringle et al., 2026).

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

More advanced hydrodynamic models such as ADvanced CIRculation (ADCIRC) (Luettich and Westerink, 2004) and Semi-implicit Cross-scale Hydroscience Integrated System Model (SCHISM) (Zhang et al., 2016) use unstructured meshes to efficiently simulate cross-scale flows and resolve detailed coastal features, and can simulate the compound dynamics of surge, rain, and river discharge. NOAA's Surge and Tide Operational Forecast Systems (STOFS) which is the U.S. gold standard for marine navigation uses ADCIRC and SCHISM for 2-D and 3-D hydrodynamic simulation of the globe and the Atlantic region respectively (Cui et al., 2024; Funakoshi, et al., 2025). However, these models are simulated on basin-scale meshes and include nonlinear terms making them more computationally expensive than SLOSH's setup in PSurge for a small target region. With current computing capacity and versions of code limited to CPUs, large-scale ensemble forecasting (hundreds of members) in operations to deliver timely results is not feasible. To address this limitation, Pringle et al. (2023) developed a machine learning (ML) based probabilistic framework for surge prediction that reduces the number of ensembles (to potentially as low as 19) to cover the uncertainty ranges typically accounted for by the PSurge workflow. In this context, we introduce Unstructured Probabilistic Surge with Machine Learning (UPSurgeML), a ML based workflow with a high resolution unstructured ocean model for probabilistic surge prediction. UPSurgeML is designed to provide rapid, cross-scale probabilistic storm surge forecast guidance during the Atlantic hurricane season. It has been tested in the deterministic form on NOAA's operational High Performance Computing (HPC) systems during the 2025 Atlantic hurricane season and is configured to provide probabilistic experimental guidance during the 2026 Atlantic hurricane season.

The rest of the manuscript is organized as follows: Section 2 describes the model structure, components, and expected outputs; Section 3 demonstrates the model's implementation using

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

three retrospective storm case studies in the Atlantic region; Section 4 presents and discusses the results for these case studies; and Section 5 summarizes the conclusions and provides closing remarks.

2. Model Description

This section provides a comprehensive overview of the code structure and its underlying components (Fig. 1). Section 2.1 describes the software architecture, followed by outlines of the various workflow components: ingesting the initial storm track input (Section 2.2); executing the simulations (Section 2.3); aggregating simulation outputs and conducting probabilistic predictions of storm surge (Section 2.4).

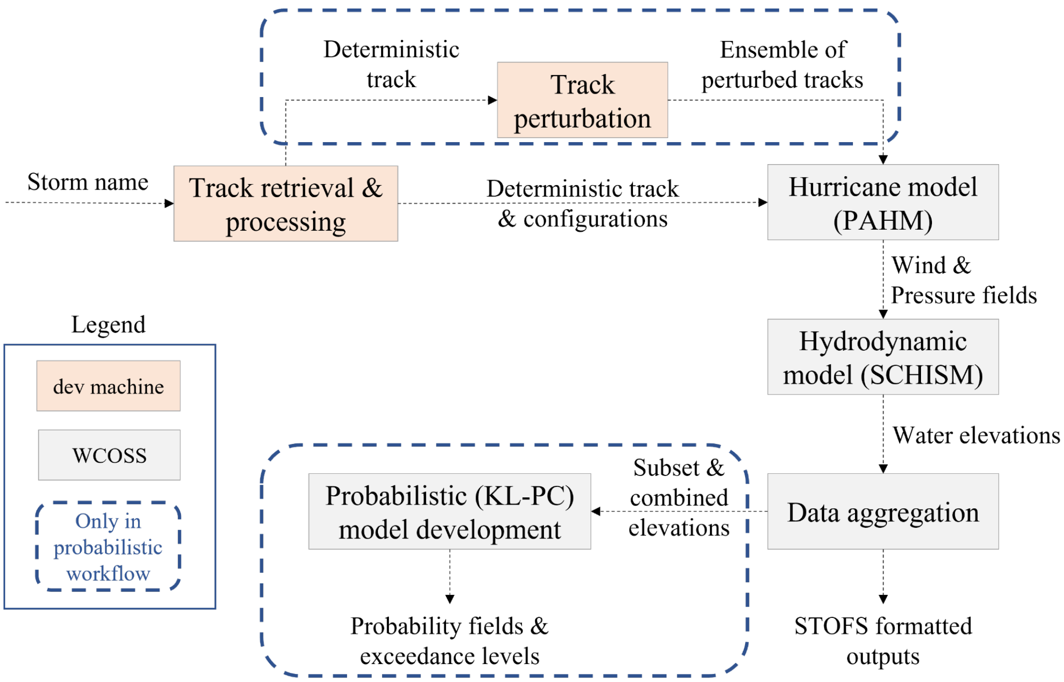

Figure 1. Overview of UPSurgeML workflow

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

2.1. Software Architecture

UPSurgeML is an open-source model written primarily in Python, while the integrated hurricane vortex and hydrodynamic model engines (see Section 2.3) are written in C and Fortran. It complies with the National Centers for Environmental Prediction (NCEP) Central Operations (NCO) standards, ensuring its suitability for implementation on NOAA's Weather and Climate Operational Supercomputing System (WCOSS) (NCEP-WCOSS, 2022). As the standard requires, UPSurgeML is compatible with the ecFlow workflow manager (ecFlow, 2026) to generate job cards and Portable Batch System (PBS) for job scheduling, maintaining a one-to-one correspondence between the ecf, jobs, and scripts directories (Fig. 2).

```
.ecf (processed to create a Job Card)
 └── jobs (J-job)
      └── scripts (ex-script, e.g., *.sh or *.py)
           ├── ush (utility scripts, e.g., *.sh or *.py)
           └── exec (compiled executables)
```

Figure 2. NCEP WCOSS implementation standard.

The package operates within the STOFS v2.1* environment available on WCOSS. Developed for on-demand execution, the workflow is initiated by a trigger mechanism adopted from the current experimental PSurge workflow. This initial step is executed on a development (dev) machine to generate a trigger package (Section 2.2), which then launches the computationally expensive operations on the HPC system (in this case, WCOSS) that includes hydrodynamic simulations forced by parametric hurricane winds (Section 2.3), data aggregation, and ML-based probabilistic evaluation of the storm surge (Section 2.4).

2.2. Trigger Package

UPSurgeML provides both deterministic and probabilistic prediction of storm surge. The process is initiated by retrieving the storm track(s) (Section 2.2.1), then through a sequence of steps, first

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

the forecast track (Section 2.2.2) and then the ensemble of perturbed tracks (Section 2.2.3) are generated on the "dev" machine. Processed tracks, along with perturbation information, and a shapefile of the region of storm impact are packaged into a tarball (that we call the "trigger package") to transfer to the HPC (in this case, WCOSS).

2.2.1. Storm Track Retrieval

UPSurgeML requires a strom track in the Automated Tropical Cyclone Forecasting System (ATCF) format (Sampson and Schrader, 2000). It works with a user-provided local files (best-track and advisory forecast tracks), and also provides two options for retrieval: 1) the online NHC ATCF dataset (called a-deck), or 2) a standalone best-track file (called b-deck, for deterministic runs only). For the first two options, the historical analysis data leading up to the simulation start time is merged with the advisory forecast track. In the third option (deterministic runs), only the best-track is utilized. By default, UPSurgeML uses the open-source StormEvents (2026) Python package to retrieve storm tracks from the online ATCF portal.

2.2.2. Process the Storm Track

Once the track is retrieved, UPSurgeML uses one of multiple parametrization options for the prediction of missing values of central pressure and RMW. For central pressure, the supported methods include, assuming persistence of the initial (0-hr forecast time) Holland *B* (Holland, 1980) parameter (Pringle et al., 2026), Courtney and Knaff (2009), or Chavas et al (2025). Based on our tests, the Courtney and Knaff (2009) method is recommended, while the persistent Holland *B* assumption, which we used in our initial beta version, is discouraged. For RMW prediction, the available methods include, persistence of the initial (0-hr forecast time) value, and the regression method developed by Penny et al. (2023), applied either with or without 24-hr moving mean smoothing. The Penny et al. (2023) method with smoothing is recommended (Pringle et al., 2026).

Table 1 summarizes the specific combinations of these parameterizations supported in UPSurgeML, alongside with their predefined configuration tags. Pringle et al. (2026) evaluated the performance of these methods on the probabilistic prediction of coastal flooding using the KL-PC surrogate framework (Pringle et al., 2023). By default, UPSurgeML utilizes the recommended NS05 configuration to predict missing central pressure and RMW values.

Table 1. UPSurgeML options for prediction of missing central pressure and the radius of maximum wind (RMW)

Tag	Central pressure prediction method	RMW prediction method
NS01	Holland (1980)	Penny et al. (2023) regression with smoothing
NS02	Holland (1980)	Persistent
NS03	Holland (1980)	Penny et al. (2023) regression without smoothing
NS04	Chavas et al (2025)	Penny et al. (2023) regression with smoothing
NS05	Courtney and Knaff (2009)	Penny et al. (2023) regression with smoothing

2.2.3. Perturb Forecasted Storm Track

In order to generate an ensemble of tracks (for probabilistic analysis), four parameters representing a storm's intensity, size, and location, including 1) cross-track location (CT), 2) along-track location (AT), 3) maximum wind speed (V_{max}), and 4) RMW are perturbed from the NHC advisory forecast. The low-discrepancy Korobov sequence (Korobov, 1959) recommended by Pringle et al. (2023) is used to sample from the 5-year historical average errors of NHC advisory forecasts (Penny and Cangialosi, 2019). In addition to V_{max} and RMW, NHC advisory forecasts contain information on the wind radii (30, 50, and 64-kt isotachs) in each storm quadrant. These are treated

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

as dependent parameters, and are perturbed based on the perturbations of the four independent parameters given the assumption that the shape parameters of the wind profile remain constant (see Pringle et al. (2026) for details). The ensemble of perturbed tracks (40 by default) are generated at this stage. Generated track(s) and associated files called "trigger package" are then transferred to WCOSS for hurricane and hydrodynamic simulations.

2.3. Simulation

Parametric Holland Model (PAHM) (Velissariou, 2021) is used for the vortex simulations. It ingests ATCF-formated hurricane tracks and generates TC surface wind and pressure forcing files for hydrodynamic simulations. UPSurgeML uses SCHISM for ocean modeling (Zhang et al., 2016). SCHISM is an open-source hydrodynamic model that supports 2D and 3D barotropic and baroclinic modeling on unstructured meshes, enabling cross-scale simulation from inland creek to open ocean.

UPSurgeML utilizes a 2D (depth-averaged) barotropic configuration. It employs STOFS-3D Atlantic v2 (Cui et al., 2024) unstructured mesh for the Atlantic region, which consists of over 2.9 million nodes and 5.6 million elements with spatial resolution ranging from approximately 5 m to 8.9 km. SCHISM simulations are initiated with 12 days of spinup run before the forecast time, followed by the main (hotstart) runs. Simulations are conducted using 240 seconds (four minutes) timesteps. Water elevations for the original track are saved hourly, while ensemble member data is recorded every three steps (12 minutes) to ensure sub-hourly extreme values are captured. To manage storage, outputs for every hour are saved in separate files, resulting in multiple files per member forecast. Simultaneously, through the co-processing step, hourly maximums are extracted from the sub-hourly output files, and are saved in a separate file for post-processing analysis.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

2.4. Post-Processing

2.4.1. Data Aggregation

For deterministic simulation of the best-track, hourly water elevations are used to generate STOFS formatted outputs (Cui et al., 2024). That includes hourly and total maximum water elevation fields (at all computational nodes), and six-minute water elevations at CO-OPS observation locations (CO-OPS, 2026).

For probabilistic simulations, first a subset of hourly maximum water elevations in the vicinity of the storm trajectory and landfall location are extracted from the SCHISM simulations of all computational nodes. Two criteria of 1) bed elevation less than 25 m, and 2) 34 kt isotach of the storm track, are used to subset the region of interest. Then aggregated outputs for the subset region are saved as STOFS formatted outputs as well as six sets of outputs compatible with PSurge time windows for probabilistic evaluations. That includes, both hourly maximum, and total maximum water elevations for 1) the entire forecast period, 2) incremental, and 3) cumulative values for user-defined (default of six-hour) intervals.

2.4.2. Probabilistic Evaluation

The Karhunen–Loève Polynomial Chaos (KL-PC) framework developed by Pringle et al. (2023) is employed to generate probabilistic predictions from the limited ensemble size. First, the spatial dimensionality of the outputs (subset of mesh nodes, on order of 10^5) are reduced through KL decomposition, obtaining approximately 10 dominant modes. Then, 3rd order PC surrogate models are developed for each of these dominant modes. Through rearrangement of the KL-PC construction, the PC coefficients (i.e., surrogate model) for each subset mesh node can then be trivially computed (see Pringle et al., (2023) for more details). Finally, Monte Carlo sampling of the KL-PC surrogates is performed to generate the probabilistic predictions. This includes 1)

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

probability of water elevation exceeding certain thresholds ranging from 1ft to 20ft, and 2) probability of water elevation exceedance for multiple thresholds ranging from 10% to 90% (with 20% intervals). These probabilistic analyses are conducted for 1) the total maximum water elevation, and user-defined (i.e., six-hour) 2) incremental maximum, and 3) cumulative maximum water elevations. Outputs are saved in netCDF files with a format similar to the PSurge GRIB files.

3. Case study

Three historical storms including Ike (2008), Sandy (2012), and Idalia (2023) are used to showcase UPSurgeML performance across different coastal regions of the Atlantic and Gulf of America (Fig. 3). Hurricane Ike was a category 4 storm that crossed over Cuba before making the landfall as a category 2 hurricane on September 13, 2008 in Galveston, Texas. Hurricane Ike was one of the largest storms in the Gulf coast region, which made landfall with maximum sustained wind of 110 mph (177 km/h), and made substantial damage on land as it moved over Texas and Louisiana (responsible for at least 49 fatalities and over $24 billion damage in the United States), which made it the third devastating storm to the United States at the time, after Hurricanes Katrina (2005) and Andrew (1997) (Berg, 2014). Hurricane Sandy was a late-season hurricane with multiple landfalls (as category 1 and 3 storms) in the Caribbean region before reaching the United States. Despite weakening, it continued to grow in size before reaching the United States and made landfall as a post-tropical cyclone (category 1 equivalent) on October 29, 2012 in New Jersey. Sandy caused at least 147 direct deaths across the Atlantic. With near $50 billion damage to the United States, it was the second most damaging storm since 1900 (Blake et al., 2013). Hurricane Idalia was a category 4 storm, and made landfall on the Florida Big Bend on August 30, 2023 as a category 3 storm. It was the third strongest storm that made landfall in the Big Bend region at the time (with

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

maximum sustained wind of 125 mph / 201 km/h), and caused 12 fatalities and over $3.6 billion damage in the United States (Cangialosi and Alaka, 2024).

For all three storms, the official advisory tracks 48 hours before landfall are used to showcase UPSurgeML for creation of probability fields and exceedance levels up to 96 hours forecast (48 hours after landfall). Section 3.1 provides a step-by-step guide for executing UPSurgeML for the Hurricane Sandy case study, and the results for all three storms are presented in Section 4.

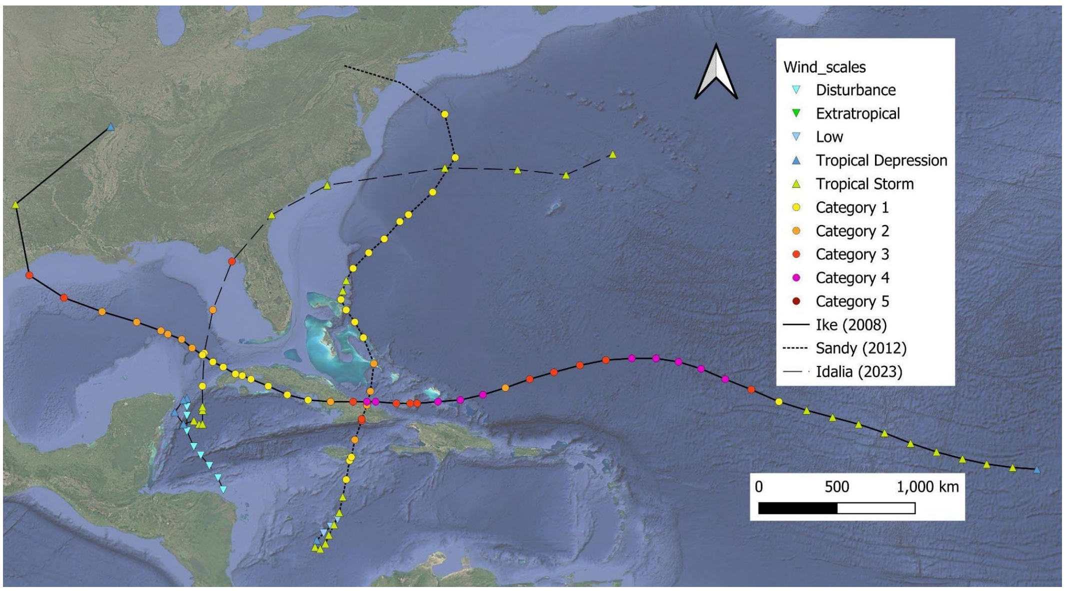

Figure 3. Trajectory and intensity of Hurricanes Ike (2008), Sandy (2012), and Idalia (2023).

3.1. Step-by-step Model Execution

The first step to run UPSurgeML is retrieving the storm track and generating the trigger package on the dev machine. The following command retrieves the NHC advisory track for hurricane Sandy (al182012) 48 hours before landfall from the web. It will generate the configuration file (with path to inputs and outputs directories), and the trigger package that contains information of 40 perturbed ensemble tracks with user-defined methods for predicting central pressure and RMW (Fig. 4).

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026


```
python ./dev/upsurgeml_preptrack.py 20121028Z06 \
~/path_to_trigger_dir/ \
~/path_to_results_dir/ \
--config-id 4 \
--rmw regression_penny_2023 \
--pc regression_courtney_knaff_2009 \
--ensemble online sandy 2012
```

Where the `config-id` is a user-defined integer for keeping track of the configurations used (i.e., auto-trigged simulations) for model setup. It should be noted that running `*.sh` files (`dev/upsurgeml_preptrack.sh`, or `dev/upsurgeml_preptrack_cron.sh`) instead of the Python code (shown above), will use default values and will have less flexibility for choosing data retrieval methods and for determining how to handle missing values.

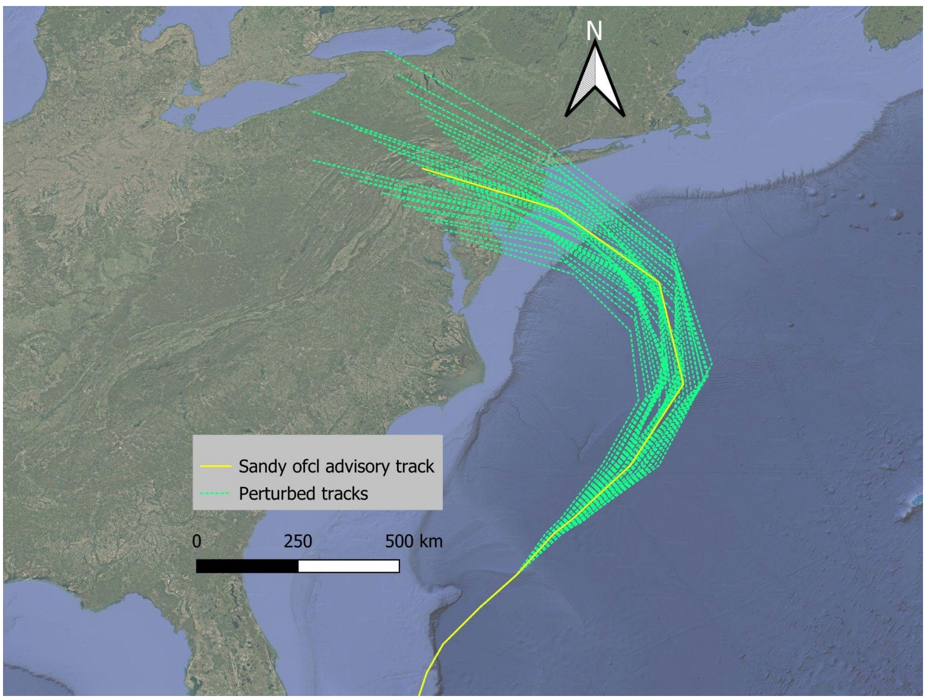

Figure 4. Hurricane Sandy official advisory track 48-hours before landfall and 40 perturbations of the track (only the CT and AT trajectory perturbations are shown here) generated by UPSurgeML

The trigger config file generated from the command above contains the following line:

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026


```
al182012 00 /path_to_tracks.dat 20121028 06 /path_to_results_storage/
```

Once the trigger and configuration files are transferred to WCOSS, the following command will run the simulations, conduct the co-processing step, generate aggregated files for probabilistic analysis, and conduct probabilistic analysis:

```
./dev/upsurgeml_cronjob.sh 4
```

Where the number at the end defines the configuration ID (explained above).

4. Results and Discussion

UPSurgeML provides probabilistic prediction of user-defined probability fields and exceedance levels for a) maximum, b) 6-hour incremental, and c) 6-hour cumulative water elevations. A few selected sets of results are showcased here. Fig. 5 shows 50% (median) and 10% exceedance (90th percentile) of maximum water elevations of the three storms with 48 hours lead time. Hurricane Ike resulted in the highest elevations near the landfall location, followed by Sandy and Idalia. The median water elevation near the landfall location of Ike exceeded 3 m, but was less than 3 m and 2.5 m in the vicinity of Sandy and Idalia landfall locations respectively. The 90th percentile (10% exceedance) of water elevations exceeded 3 m for all storms. Comparison of exceedance levels with maximum water elevations at the CO-OPS stations revealed that the 90th percentile well represents maximum recorded water elevations especially in the vicinity of landfall locations.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

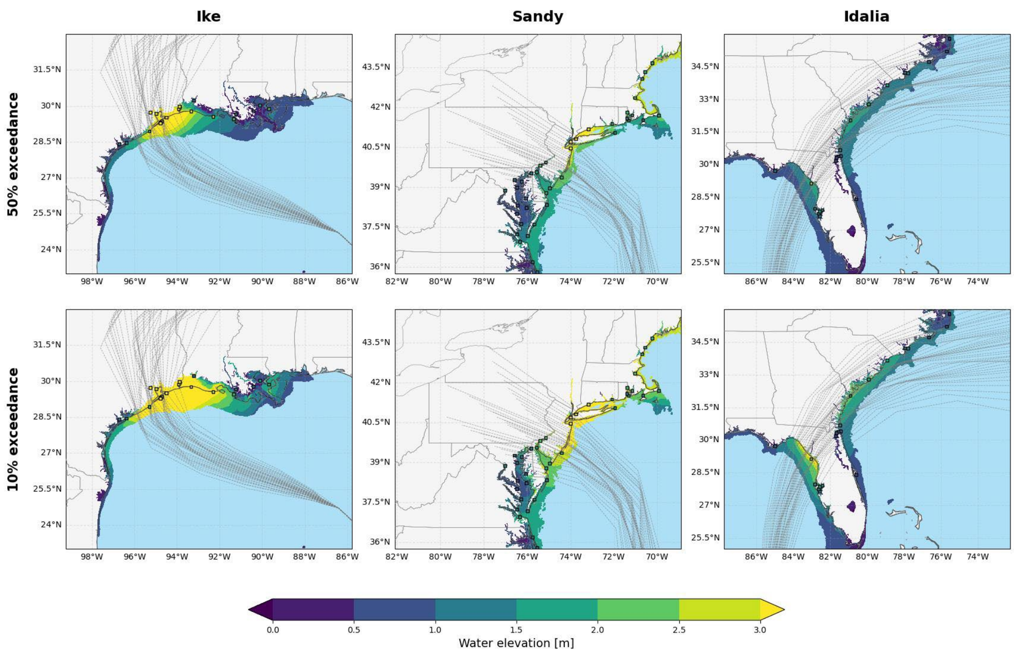

Figure 5. 50% (median), and 10% exceedance probability fields of the maximum water elevation for Hurricanes Ike, Sandy, and Idalia. Ensembles of hurricane tracks are shown in gray. Maximum water elevations at CO-OPS stations are shown with black squares.

The probability of maximum water elevation exceeding 3 ft (~1 m) and 6 ft (~2 m) is shown in Fig. 6. For Hurricane Ike, the region exceeding 3 ft threshold is contained to the vicinity of the hurricane track, but for Sandy and Idalia, a wider region across the Atlantic coast exceeded the 3ft threshold, ranging from the East Coast of Florida to the North East New England region. This is mostly a consequence of the tidal environment – microtidal in the Gulf of America and mesotidal along most of the Atlantic Coast. This is why the 6 ft exceedance impact region is notably smaller, especially for Idalia. Though indeed, Sandy's impact region by this threshold is the widest, given the large-scale nature of the storm.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

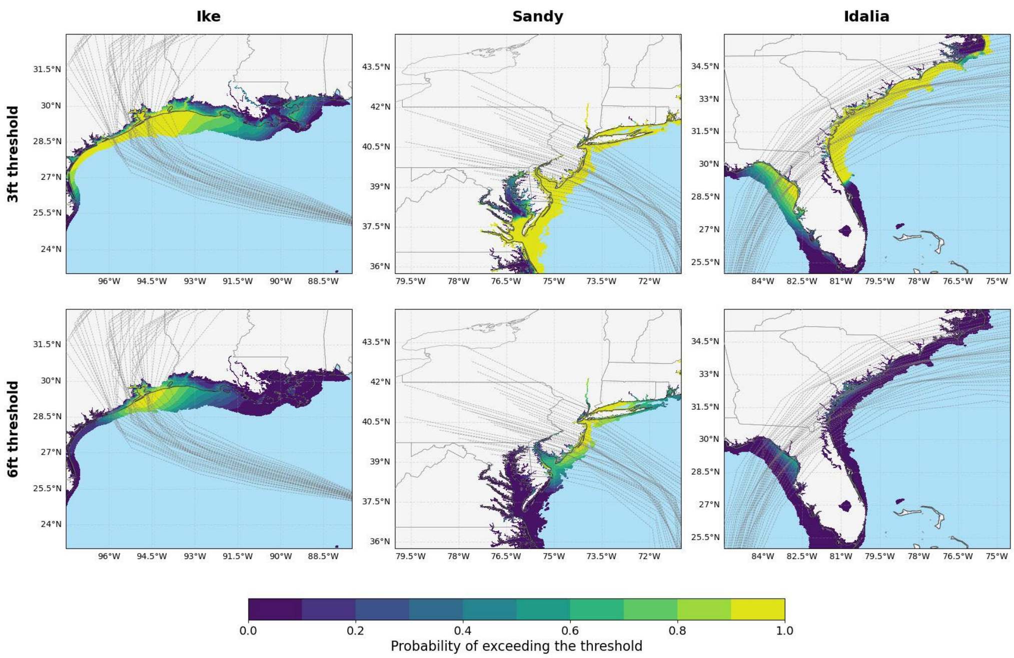

Figure 6. Probability of maximum water elevations exceeding 3 ft and 6 ft thresholds for Hurricanes Ike, Sandy, and Idalia. Ensembles of hurricane tracks are shown in gray.

In addition to probabilistic prediction of maximum water elevations, UPSurgeML provides incremental and cumulative assessments for user-defined (here 6-hour) intervals during the (96 hours) simulation period. Fig. 7 shows the incremental probability of exceeding 6ft threshold for 6-hours before landfall, the landfall time, 6-hours and 12-hours afterward. For Ike, the probability of exceeding 6 ft in the vicinity of the storm track is low (<0.5) up to 6 hours before and 12 hours after the landfall, but very high probabilities (~1.0) close to or during the landfall time. Probabilistic evaluation of maximum water elevations are useful for overall assessment of storm surge, but incremental and cumulative results provide more details about the timing of extreme values, which is crucial for timely emergency responses, and the issuance of early warning alerts.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

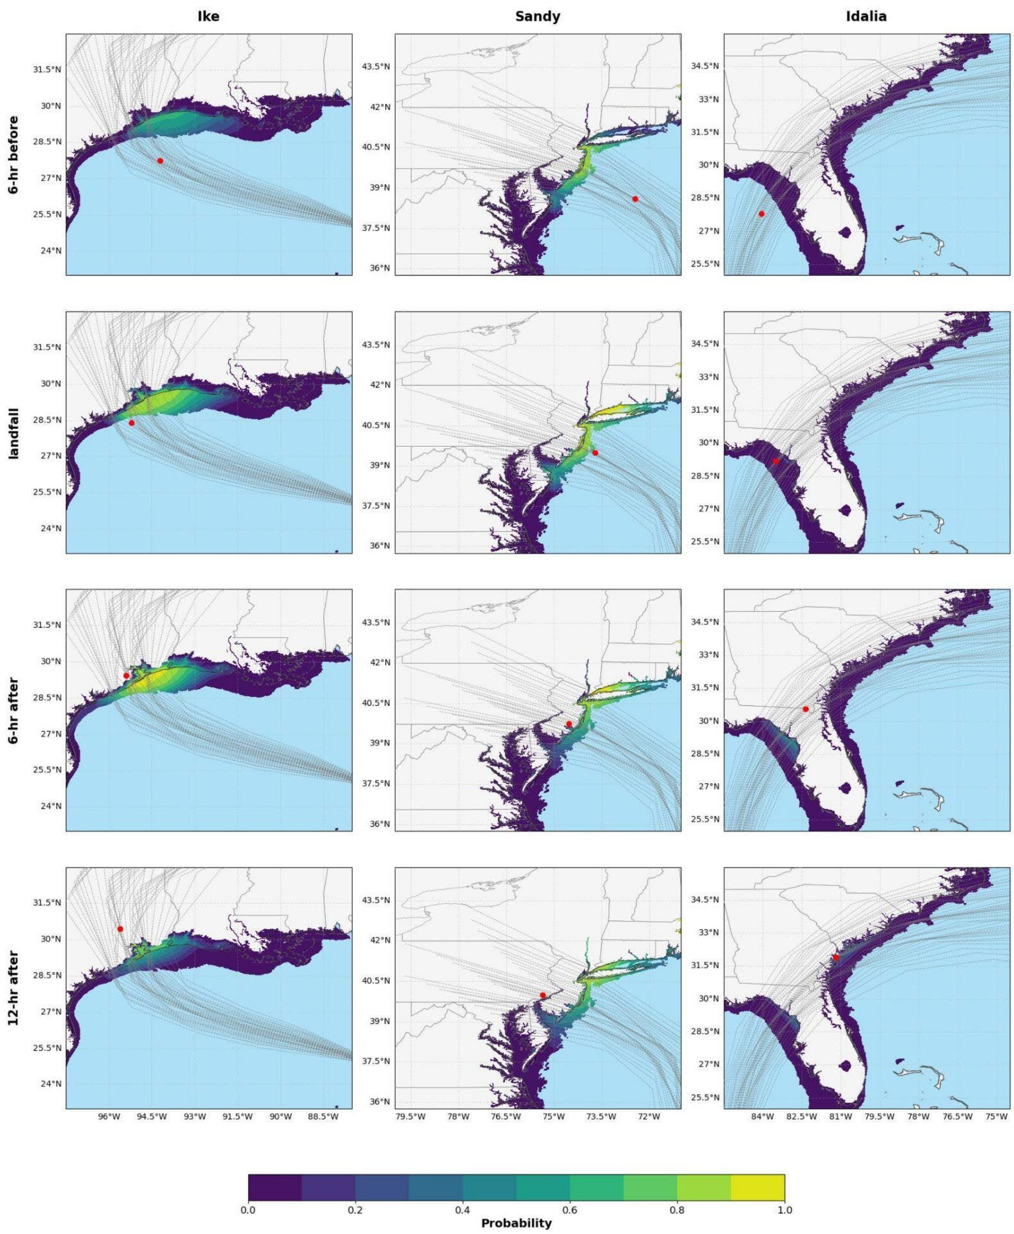

Figure 7. Probability of water elevations exceeding 6ft during four consecutive 6-hour intervals including: 6-hour before landfall, landfall time, 6-hour and 12-hour after landfall for Hurricanes Ike, Sandy, and

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

Idalia. Ensembles of hurricane tracks are shown in gray. Red points show the center of the official advisory track at the time of plot.

5. Conclusion

This study presented the UPSurgeML (v1.0), a semi-operational workflow for probabilistic assessment of TC-driven storm surge. It's a Python-driven and NCO compliant workflow which makes it suitable for implementation on NOAA's HPCs (i.e., WCOSS) during the hurricane season. The deterministic form of the workflow was tested during the 2025 Atlantic hurricane season and the probabilistic form will be tested during the upcoming 2026 Atlantic hurricane season. The process is initiated with generation of a trigger package on the dev machine, which contains hurricane track(s) and configurations for parametric wind model and hydrodynamic simulations. The wind model (PAHM) and hydrodynamic simulations (with SCHISM) are then conducted on WCOSS. A machine learning surrogate model framework is used to make probabilistic predictions and output exceedance water level and probability fields. The workflow performance was showcased for three historic storms across the Atlantic region including Ike (2008), Sandy (2012), and Idalia (2023). UPSurgeML uses a similar triggering process as the operational PSurge and generates similar probabilistic outputs, but unlike PSurge, it uses an unstructured grid that enables cross-scale simulation from inland creek to open ocean.. The machine learning framework of UPSurgeML also significantly reduces the number of ensembles compared to PSurge for probabilistic prediction (~400-1200 to 40).

The current workflow only uses atmospheric forcing for surge modeling. Integration of inland hydrology and wave setup in the future versions could improve surge prediction especially for wet storms (e.g., Hurricane Harvey 2017; Huang et al., 2021) and regions where wave setup effects are significant (e.g., Puerto Rico; Joyce et al., 2019). Another improvement to reduce

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

computational cost can be generating storm specific mesh subsets with high-resolution around the impact region (Cassalho et al., 2026) instead of using the STOFS-3D v2 mesh, which refines the entire U.S. Atlantic and Gulf coastline with high-resolution.

Code and Data Availability

The current version of UPSurgeML is available on NOAA's Office of Coast Survey (OCS) modeling GitHub website: https://github.com/noaa-ocs-modeling/UPSurgeML. The version of the model used for this manuscript, along with input and output files and the additional scripts used to generate plots are available at https://doi.org/10.5281/zenodo.20801237 (Daneshvar et al. (2026).

Author Contributions

FD: conceptualization, methodology, code development, visualization, model validation, writing (original draft); SMa: project management, conceptualization ,methodology, code development, data curation, workflow development and implementation, workflow validation, writing (review and editing); WJP: conceptualization, methodology, code development, model validation, writing (review and editing); PV: conceptualization, methodology, model validation, writing (review and editing), ZY: data curation, writing (review and editing); GS: data curation, writing (review and editing); EM: conceptualization, principal investigator; SMo: conceptualization, principal investigator, writing (review and editing).

Competing Interests

The authors declare that they have no conflict of interest.

Disclaimer

The scientific results and conclusions, as well as any views or opinions expressed herein, are those of the author(s) and do not necessarily reflect those of NOAA or the Department of Commerce.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

Acknowledgements

The authors would like to acknowledge Khachik Sargsyan who contributed to the conceptual theory and development of the KL-PC framework, and Zachary Burnett who contributed to initial development of the ensemble perturbation code, a fundamental component of UPSurgeML. We also would like to thank the National Hurricane Center's Andrew Penny and Laura Alaka for their valuable feedback and peer evaluation of UPSurgeML. The workflow described in this manuscript is designed to work and also tested on NOAA's High Performance Computing (HPC). Simulations used in this manuscript were also conducted on the same HPC.

Financial Support

This work was supported by funding from NOAA through Disaster Relief Supplemental Appropriations Act, 2022 and from H.R.3684 Infrastructure Investment and Jobs Act (IIJA) projects. WJP is supported by NOAA grant NA23NOS4730238 through a Strategic Partnership Project agreement (A23089) between Virginia Institute of Marine Science and Argonne National Laboratory under U.S. Department of Energy contract DE-AC02-06CH11357

References

Berg, R.: Tropical Cyclone Report: Hurricane Ike (AL092008), National Hurricane Center, 2014.

Blake, E. S., Kimberlain, T. B., Berg, R. J., Cangialosi, J. P., and Beven II, J. L.: Tropical Cyclone Report: Hurricane Sandy (AL182012), National Hurricane Center, 2013.

Cangialosi, J. P. and Alaka, L.: Tropical Cyclone Report: Hurricane Idalia (AL102023), National Hurricane Center, 2023.

Cassalho, F., Mani, S., Moghimi, S., Ye, F., and Zhang, Y.: OCSMesh and an automated creek-to-ocean mesh generation workflow. Ocean Modelling, 203, 102774, https://doi.org/10.1016/j.ocemod.2026.102774, 2026.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

Chavas, D. R., Knaff, J. A., and Klotzbach, P.: A simple model for predicting tropical cyclone minimum central pressure from intensity and size, Weather and Forecasting, 40, 333–346, https://doi.org/10.1175/WAF-D-24-0031.1, 2025.

Courtney, J. and Knaff, J. A.: Adapting the Knaff and Zehr wind-pressure relationship for operational use in Tropical Cyclone Warning Centres, Australian Meteorological and Oceanographic Journal, 58, 167–179, https://doi.org/10.1071/ES09021, 2009.

Cui, L., Ye, F., Zhang, Y. J., Yu, H., Wang, Z., Moghimi, S., Seroka, G., Riley, J., Pe'Eri, S., Mani, S., and Myers, E.: Total water level prediction at continental scale: Coastal ocean, Ocean Modelling, 192, 102451, https://doi.org/10.1016/j.ocemod.2024.102451, 2024.

Daneshvar, F., Mani, S., Pringle, W. J., Velissariou, P., Sargsyan, K., Yang, Z., Seroka, G., Myers, E., and Moghimi, S.: UPSurgeML (v1.0) - A Machine Learning Workflow for Probabilistic Ensemble Forecast of Tropical Cyclone Storm Surge: Data and Code, Zenodo [code and data set], https://doi.org/10.5281/zenodo.20801237, 2026.

Daneshvar, F., Mani, S., Pringle, W. J., Moghimi, S., and Myers, E.: Description of a Python script to implement a probabilistic surrogate model for predicting tropical cyclone driven flood depth and extent, NOAA Technical Memorandum NOS CS 59, https://doi.org/10.25923/wyvn-4405, 2024.

ecFlow: ecFlow documentation, available at: https://ecflow.readthedocs.io/en/latest/ (last access: 18 May 2026), 2026.

Funakoshi, Y., Seroka, G., Moghimi, S., Myers, E., Alipour, A., Haddad, J., Mani, S., Westerink, J., Cerrone, A., Tejaswi, A., Blakely, C., and Wood, D.: NOAA's Surge and Tide Operational Forecast System 2D Global (STOFS-2D-Global) Version 2, https://doi.org/10.25923/y34r-z463, 2025.

Glahn, B., Taylor, A., Kurkowski, N., and Shaffer, W. A.: The role of the SLOSH model in National Weather Service storm surge forecasting, National Weather Digest, 33, 2009.

Holland, G. J.: An analytic model of the wind and pressure profiles in hurricanes, Monthly Weather Review 108, 1212–1218, https://doi.org/10.1175/1520-0493(1980)108<1212:AAMOTW>2.0.CO;2, 1980.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

Jelesnianski, C. P., Chen, J., and Shaffer, W. A.: SLOSH: Sea, Lake, and Overland Surges from Hurricanes, NOAA Technical Report NWS 48, 1992.

Joyce, B. R., Gonzalez-Lopez, J., van der Westhuysen, A. J., Yang, D., Pringle, W. J., Westerink, J. J., and Cox, A. T.: U.S. IOOS coastal and ocean modeling testbed: Hurricane-induced winds, waves and surge for deep-ocean, reef fringed islands in the Caribbean, Journal of Geophysical Research: Oceans, 124, 2876–2907, https://doi.org/10.1029/2018JC014687, 2019.

Kerr, P. C., Donahue, A. S., Westerink, J. J., Luettich, R. A., Zheng, L. Y., Weisberg, R. H., Huang, Y., Wang, H. V., Teng, Y., Forrest, D. R., and Roland, A.: U.S. IOOS coastal and ocean modeling testbed: Inter-model evaluation of tides, waves, and hurricane surge in the Gulf of Mexico, Journal of Geophysical Research: Oceans, 118, 5129–5172, https://doi.org/10.1002/jgrc.20376, 2013.

Korobov, N. M.: Approximate evaluation of repeated integrals, Dokl. Akad. Nauk SSSR, 124, 1207–1210, 1959.

Luettich, R. A. and Westerink, J. J.: Formulation and numerical implementation of the 2D/3D ADCIRC finite element model version 44.XX, University of North Carolina at Chapel Hill and University of Notre Dame Technical Report, 74 pp., 2004.

NCEP-WCOSS: NCEP WCOSS Implementation Standards, available at: https://nws-hpc-standards.readthedocs.io/en/latest/ (last access: 22 April 2026), 2022.

NOAA Center for Operational Oceanographic Products and Services (CO-OPS): NOAA Tides and Currents, available at: https://tidesandcurrents.noaa.gov/ (last access: 20 June 2026), 2026.

PAHM: Parametric Hurricane Model, GitHub [code], available at: https://github.com/noaa-ocs-modeling/PaHM (last access: 23 April 2026), 2022.

Penny, A. B., Alaka, L., Taylor, A. A., Booth, W., DeMaria, M., Fritz, C., and Rhome, J.: 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, 38, 2461–2479, https://doi.org/10.1175/WAF-D-22-0209.1, 2023.

Penny, A. and Cangialosi, J. P.: P-surge error update, NOAA Technical Report, 13 pp., 2019.

https://doi.org/10.5194/egusphere-2026-4586
Preprint. Discussion started: 15 September 2026

Pringle, W. J., Burnett, Z., Sargsyan, K., Moghimi, S., and Myers, E.: Efficient probabilistic prediction and uncertainty quantification of tropical cyclone–driven storm tides and inundation, Artificial Intelligence for Earth Systems, 2, e220040, https://doi.org/10.1175/AIES-D-22-0040.1, 2023.

Pringle, W. J., Mani, S., Daneshvar, F., Brannan, A., Velissariou, P., Alaka, L., Zhang, J., Moghimi, S., and Fritz, C.: Cross-scale Probabilistic Storm Surge and Tide Forecasts: Impacts of Tropical Cyclone Wind Field Structure, ESS Open Archive [preprint], https://doi.org/10.22541/essoar.15004678/v1, 2026.

PSurge: MDL PSurge, available at: https://vlab.noaa.gov/web/mdl/psurge (last access: 30 April 2026), 2026.

Sampson, C. R. and Schrader, A. J.: The automated tropical cyclone forecasting system (version 3.2), Bulletin of the American Meteorological Society, 81, 1231–1240, 2000.

StormEvents: StormEvents, GitHub [code], available at: https://github.com/oceanmodeling/StormEvents (last access: 23 April 2026), 2026.

Velissariou, P.: Development of a Flexible Parametric Hurricane Modeling System (PaHM) for Coastal Coupled Applications, GitHub [code], available at: https://github.com/noaa-ocs-modeling/PaHM (last access: 27 April 2026), 2021.

Zhang, K., Xiao, C., and Shen, J.: Comparison of the CEST and SLOSH Models for Storm Surge Flooding, Journal of Coastal Research, 24, 489–499, https://doi.org/10.2112/06-0709.1, 2008.

Zhang, Y., Ye, F., Stanev, E. V., and Grashorn, S.: Seamless cross-scale modeling with SCHISM, Ocean Modelling, 102, 64–81, https://doi.org/10.1016/j.ocemod.2016.05.002, 2016.