

Correcting Physics-Based Global Tide and Storm Water Level Forecasts with the Temporal Fusion Transformer

A.R. Cerrone^{1,2}, L.G. Westerink^{1,3}, G. Ling^{1,4}, C.P. Blakely¹, D. Wirasaet¹, C.
Dawson², and J.J. Westerink¹

¹Department of Civil & Environmental Engineering & Earth Sciences, University of Notre Dame, Notre
Dame, IN, USA 46556

²Oden Institute for Computational Engineering & Sciences, University of Texas at Austin, Austin, TX,
USA 78712

³Department of Computer Science, University of Texas at Austin, Austin, TX, USA 78712

⁴International Research Institute of Disaster Science, Tohoku University, Sendai, Miyagi, Japan

Key Points:

- Global water level forecasting models are subject to several sources of epistemic uncertainty
- Machine-learning-based transformer models can rapidly correct for model discrepancy stemming from systematic bias and epistemic uncertainty
- Coupling transformer models to high-fidelity global water-level models is highly effective when supplemented with physics-based covariates

Corresponding author: Albert Cerrone, acerrone@nd.edu

Abstract

Global and coastal ocean surface water elevation prediction skill has advanced considerably with improved algorithms, more refined discretizations and high-performance parallel computing. Model skill is tied to mesh resolution, the accuracy of specified bathymetry/topography, dissipation parameterizations, air-sea drag formulations, and the fidelity of forcing functions. Wind forcing skill can be particularly prone to errors, especially at the land-ocean interface. The resulting biases and errors can be addressed holistically with a machine-learning (ML) approach. Herein, we weakly couple the Temporal Fusion Transformer to the National Oceanic and Atmospheric Administration’s (NOAA) Storm and Tide Operational Forecast System (STOFS 2D Global) to improve its forecasting skill throughout a 7-day horizon. We demonstrate the transformer’s ability to enrich the hydrodynamic model’s output at 228 observed water level stations operated by NOAA’s National Ocean Service. We conclude that the transformer is a rapid way to correct STOFS 2D Global forecasted water levels provided that sufficient covariates are supplied. For stations in wind-dominant areas, we demonstrate that including past and future wind-speed covariates make for a more skillful forecast. In general, while the transformer renders consistent corrections at both tidally and wind-dominant stations, it does so most aggressively at tidally-dominant stations. We show notable improvements in Alaska and the Atlantic and Pacific seaboard of the United States. We evaluate several transformers instantiated with different hyperparameters, covariates, and training data to provide guidance on how to enhance performance.

Plain Language Summary

Forecasted water levels in coastal regions are used to predict flooding risk, currents that impact navigation, the transport of nutrients and pollutants, and the conditions that make coastal waters favorable for fisheries. Computational models are typically exercised to render water-level forecasts as far as seven days into the future. While the physics underlying these models is relatively well understood, the data flowing into the models can be subject to uncertainty. For example, water depth is a sensitive model parameter, yet even the best LiDar-derived data sets may have significant errors. Model and meteorological wind speed data is another parameter prone to bias and uncertainty. In some cases, these models lack the ability to resolve oceanic and hydrological processes, phenomena that contribute to coastal water levels. Herein, we propose a fairly simple machine-learning-based strategy to improve the predictive capacity of these models in the presence of these limitations. In particular, we couple the ocean hydrodynamic model ADCIRC to a transformer to improve ADCIRC’s predictions. We show that the two models individually have larger error bands than when they are combined.

1 Introduction

Water levels in the coastal ocean and its adjacent floodplain impact navigation, water quality, fisheries, and livelihoods in coastal communities. Water motion is driven by highly predictable gravitational forces of the moon and sun acting on the earth’s ocean water; by more chaotic forces induced by wind, atmospheric pressure, wind waves, rainfall and regional hydrology; and forces derived from the ocean’s thermohaline structure with corresponding ocean current systems and ice packs. The governing physics was pioneered by Laplace (1776) and later expanded by Barré de Saint Venant and Boussinesq (Hager et al., 2019) leading to the depth integrated shallow water equations (Hervouet, 2007). More detailed three dimensional forms of the conservation laws were derived from the Navier-Stokes equations, and are typically subject to the assumption of a hydrostatic pressure approximation and the Boussinesq approximation for density (Roelvink & van Banning, 1995; Haidvogel et al., 2000; Hervouet, 2007). Given sufficient resolution of the geometric and bathymetric intricacies of the coastal ocean and especially of the adjacent

69 floodplain with its narrow tidal inlets, estuaries, and dendritic channel networks and rivers,
 70 the two dimensional shallow water equations generally simulate coastal water levels re-
 71 markably well. This has especially been the case as open water boundary condition spec-
 72 ification uncertainty has been eliminated by evolving from regional to global domain mod-
 73 els, as was initially envisioned by Laplace. Additionally, closure relationships focus on
 74 parameterizing bottom boundary layer dissipation, internal tide generation and subse-
 75 quent dissipation, viscous turbulent dissipation, and air-sea and air-ice-sea momentum
 76 transfer, which are all largely empirically derived from observed data.

77 Global total water level models include Deltares' Delft3D based Global Tide and
 78 Surge Model (GTSM) (Verlaan et al., 2015; De Kleermaeker et al., 2017), Environment
 79 and Climate Change Canada's NEMO-based model (Wang et al., 2021; Wang & Bernier,
 80 2023), and NOAA's ADCIRC-based unstructured mesh Storm and Tide Operational Fore-
 81 cast System 2D Global (STOFS 2D Global) model (Seroka et al., 2023). STOFS 2D Global,
 82 the model applied in this study, is a barotropic implementation of the shallow water equa-
 83 tions with resolution ranging from 80 m within intricate inlets to 25 km across the abyssal
 84 plains of the deep ocean (Pringle et al., 2021; Blakely et al., 2022). The model leads to
 85 excellent predictability of tides with the global observed-to-modeled M_2 tide compared
 86 at 236 deep water stations resulting in a coefficient of determination, R^2 , equal to .985,
 87 a mean amplitude error of 2.4 cm, and a normalized root mean square error (NRMSE)
 88 equal to 0.075. At 449 shelf stations, the model yields an R^2 equal to .984, a mean am-
 89 plitude error of 4.3 cm, and a NRMSE equal to 0.084. STOFS 2D Global is to date the
 90 most accurate published non-data assimilated model with respect to tides (Stammer et
 91 al., 2014; Blakely et al., 2022). A multiyear hindcast (see Section 2) indicates that sur-
 92 face water elevation can be predicted at 213 U.S. NOAA National Ocean Service (NOS)
 93 water level stations with an R^2 equal to 0.94, an average absolute error equal to 7.3 cm,
 94 and a NRMSE equal to 0.21.

95 Model accuracy is influenced by factors such as geometric representation and mesh
 96 resolution, topo-bathymetry, and specified values for the parameterized dissipation terms.
 97 Additionally, the underlying physical processes that are incorporated into STOFS 2D
 98 Global are subject to formulations errors. Some salient processes may be missing alto-
 99 gether. The biggest source of epistemic uncertainty associated with STOFS 2D Global
 100 is its meteorological forcing. Available wind fields, which are essential to forecasting ac-
 101 curate water levels, are typically spatially and temporally coarse, often leading to sig-
 102 nificant wind and associated water level prediction errors over shallow inland waterbod-
 103 ies. Moreover, model winds tend to be muted for intense tropical cyclones which result
 104 in muted storm surge forecasts. Collectively, these varied sources of model discrepancy
 105 motivate an improvement capability that scrutinizes simultaneously STOFS 2D Global
 106 output, its input (e.g. forcings), and past observations.

107 Generally, Bayesian methods have been adopted to improve model accuracy amidst
 108 quantifiable uncertainty. Typically, for spatially and/or temporally invariant applications,
 109 algorithms like the Extended Kalman filter (Holland, 2020) and Particle filter (Ristic et
 110 al., 2003; Li et al., 2017; Elfring et al., 2021) have been leveraged to render these improve-
 111 ments. While they have been adopted for spatiotemporal domains (Butler et al., 2012,
 112 2015; Rougier et al., 2023) and will invariably be scaled for operations in the future, due
 113 to computational complexity and expense, they cannot be deployed today for rapid hy-
 114 drodynamic model improvement. We propose to carry out model improvement by means
 115 of machine learning with a transformer model. Rather than operating on the model di-
 116 rectly, as is commonly done in a Bayesian framework, the transformer would enrich the
 117 model outputs themselves (see schematic diagrams in Figure 1). By moving to this machine-
 118 learning-based approach, although we lose some accuracy that would come from a tight
 119 coupling required by a traditional Bayesian scheme, we gain processing speed. We de-
 120 velop and demonstrate this approach in computing corrected water levels for the STOFS
 121 2D Global model.

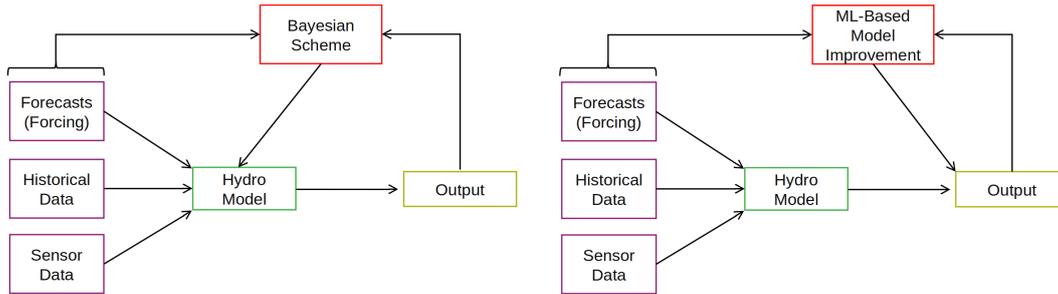


Figure 1. Process maps incorporating a Bayesian scheme (left) and ML (right) for model improvement.

122 The machine learning (ML) model considered herein is the temporal fusion trans-
 123 former (TFT) (Lim et al., 2021). The TFT is a transformer-based model (Vaswani et
 124 al., 2017) that uses self-attention to learn long-term dependencies. The transformer is
 125 a compelling option for hydrodynamic applications because some coastal regions expe-
 126 rience predictable, yet “longer term” water elevation patterns that recurrent neural net-
 127 works (RNNs) or long short term memory (LSTM) cells alone could not otherwise cap-
 128 ture. Granted, in the present application, the transformer would be compelled to scruti-
 129 nize patterns in model error space, but like the surface water elevations themselves, these
 130 error signals are also expected to be repeated. Moreover, the TFT can accommodate static,
 131 past, and future covariates. These covariates enable the TFT to consider simultaneously
 132 the relationship between observed and predicted water levels and other relevant quan-
 133 tities including wind speed and location. In essence, referring back to Figure 1, these co-
 134 variates are the data proceeding from model inputs and outputs. Consequently, as the
 135 transformer scrutinizes patterns in model error space, it also attends to model inputs and
 136 outputs, drawing parallels to a traditional Bayesian scheme.

137 ML models have featured quite extensively in coastal water level modeling. For ex-
 138 ample, in an early application, de Oliveira et al. (2009) considered a multilayer percep-
 139 tron (MLP) to predict storm surge at a single station in Southeast Brazil and demon-
 140 strated reasonable performance out to 24-hour forecast horizons. Later, Ayyad et al. (2022)
 141 assessed seven ML models to predict peak storm surge height caused by tropical cyclones
 142 in the New York Metropolitan Area. They trained their models against output from the
 143 model considered in this study, ADCIRC. They determined that of the seven ML mod-
 144 els, a support vector regressor (SVR) and an ensemble of decision trees with adaptive
 145 boosting were the most performant. Xie et al. (2023) used a convolutional neural net-
 146 work (CNN) to ingest two-dimensional wind forcing for single-site water level eleva-
 147 tion forecasting in Southeast China. They demonstrated strong performance out to three-
 148 day forecasting horizons provided that a 24-hour water-level input is supplied. Tiggeloven
 149 et al. (2021) considered an ensemble of different models which included the coupled CNN
 150 and LSTM (ConvLSTM) to predict surge levels globally at more than 700 tidal stations.
 151 They also provided meteorological data as covariates. They demonstrated that the LSTM
 152 outperformed the other models, but that the CNN, provided that it was instantiated with
 153 a sufficient number of hidden layers, had the potential to outperform the lot despite its
 154 increased compute time. Additionally, they showed that their models generally performed
 155 better for higher-latitude tidal stations. Most recently, Pachev et al. (2023) demon-
 156 strated a two-tiered location-agnostic approach to predict peak storm surge from tropical cy-
 157 clones. First, they exercised a classifier to identify inundated points. Thereafter, they
 158 ran a neural network (trained with boosting) to predict the level of inundation.

159 In this paper, we deviate our approach from these coastal water level studies. While
 160 we predict observed water levels, we do so by weakly coupling a numerical model (viz.

161 STOFS 2D Global solved by ADCIRC) to the TFT. We show that compared to ADCIRC
 162 alone, this coupling results in improved station-based water level forecasts out to 7-days.
 163 This type of coupling is not without precedent, especially in atmospheric modeling. For
 164 example, Bonavita and Laloyaux (2020) demonstrated that a MLP can extend the ca-
 165 pability of the weak-constraint formulation of the 4D-Var data assimilation framework
 166 to the troposphere. Zampieri et al. (2023) used a fully-connected MLP trained on ob-
 167 servations to reduce a temperature bias over Arctic sea ice in reanalysis products. In ocean
 168 modeling, Bolton and Zanna (2019) trained CNNs on output from a quasi-geostrophic
 169 ocean model to predict eddy momentum forcing and determined that training on tur-
 170 bulent regions leads to enhanced CNN extendability. More relevant to the present dis-
 171 cussion, they conjecture that CNNs can be coupled to sparse interpolated observational
 172 data to render accurate predictions of large-scale flow in the presence of turbulence.

173 In this study, we train the TFT to correct surface water elevation, η , predictions
 174 made by ADCIRC-based STOFS 2D Global at 228 NOAA water level stations in the north-
 175 ern half of the Western Hemisphere. We use a three-year STOFS 2D Global hindcast
 176 and 6-minute resolution NOAA water level observations to instantiate the TFTs con-
 177 sidered herein. In Section 2, we summarize STOFS 2D Global and the aforementioned
 178 three-year hindcast. Additionally, we detail the ML framework considered herein includ-
 179 ing model architecture, data processing, and training. Thereafter, in Section 3, we de-
 180 tail its performance and probe its sensitivity to training set size, various hyperparam-
 181 eters, covariates, and regionality. Finally, in Section 4, we discuss how this framework
 182 could be extended to accommodate extreme weather events and be adapted for both prob-
 183 abilistic and off-station STOFS 2D Global correction.

184 2 Materials and Methods

185 2.1 STOFS 2D Global

186 2.1.1 Overview

187 STOFS 2D Global, NOAA’s operational global water level model, is based on an
 188 optimized unstructured mesh of the global ocean with the highest resolution focused on
 189 U.S. coastal waters and floodplains. The model applies 25-km resolution across deep ocean
 190 abyssal plains, down to 2.5-km resolution across steep ocean topography, and refines coasts
 191 down to 2.5 km globally (Blakely et al., 2022). U.S. inland waters and coastal features
 192 are resolved down to 80 to 125 m. This includes a large extent of the floodplains, coastal
 193 estuaries, inland channels, and levee systems of the contiguous United States, Puerto Rico,
 194 Alaska, Hawaii, and Micronesia. Figures 2 and 3 depict mesh resolution and bathymetry
 195 in the global model. It is noteworthy that smaller elements discretize mid-ocean ridges,
 196 shelf breaks, and submerged mountain chains in order to improve internal tide dissipa-
 197 tion models (Pringle et al., 2021; Blakely et al., 2022). Global and regional bathymet-
 198 ric datasets were applied including GEBCO2020 (IHO-UNESCO, 2020), RTopo-2 (Schaffer
 199 et al., 2016), Canadian CHS-NONNA100 (Fisheries & Canada, 2023), ntheta100 Grid
 200 (Beaman, 2018), and the Allen Coral Atlas (Lyons et al., 2022). In US water, various
 201 regional bathymetric data sources were used from the United States Geological Survey
 202 (USGS), NOAA, U.S. Army Corps of Engineers, and Northeast Ocean Data. The melded
 203 bathymetry is mesh scale averaged at nodes, Figure 3.

204 The modeling system is forced with tidal potential functions, associated self attrac-
 205 tion and load terms (via FES2014), winds, atmospheric pressure, and sea ice from NOAA’s
 206 Global Forecast System (NOAA-EMC, 2023). Friction force parameterizations for bot-
 207 tom boundary layer dissipation, internal tide generation/dissipation, air-sea wind drag,
 208 and air-ice-sea drag are based on previous regional studies and were optimized for STOFS
 209 2D Global (Dietrich et al., 2011; Chen et al., 2013; Hope et al., 2013; Kerr, Donahue,
 210 et al., 2013; Kerr, Martyr, et al., 2013; Pringle et al., 2018, 2021; Blakely et al., 2022).

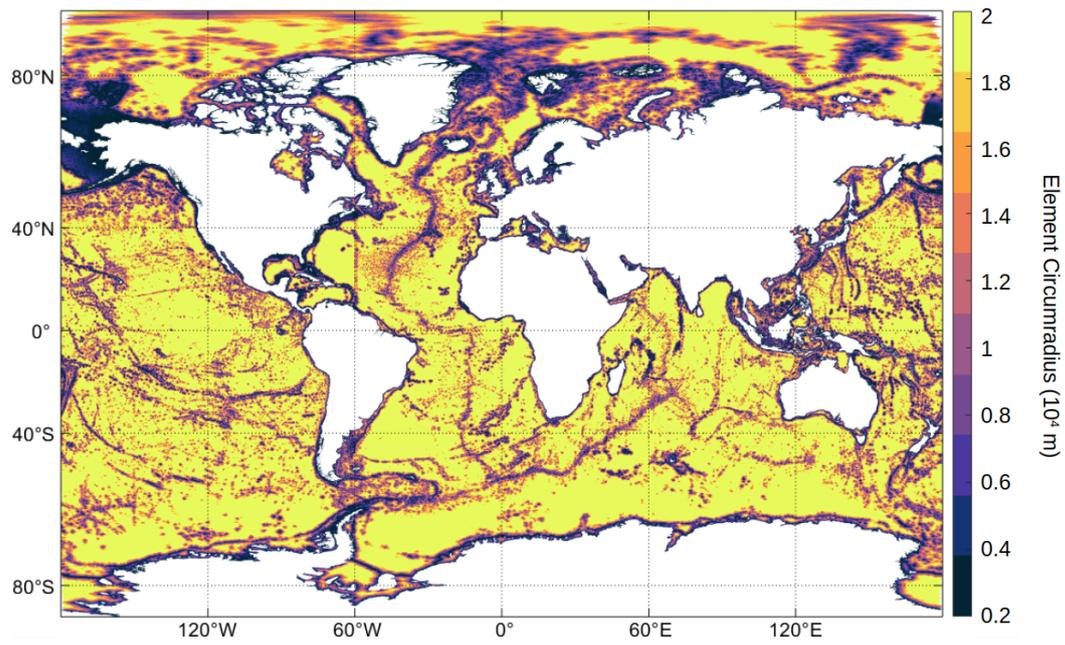


Figure 2. STOFS 2D Global mesh resolution. This mesh consists of 12,784,991 nodes and 24,875,313 elements.

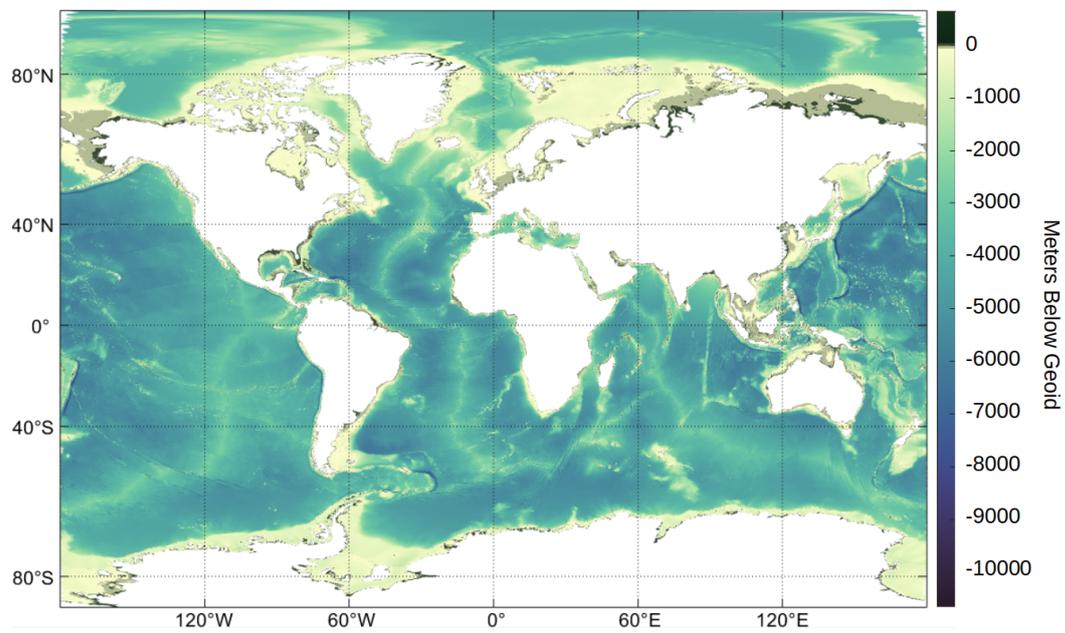


Figure 3. STOFS 2D Global bathymetry melded from various global and regional sources.

211 NOAA’s National Centers for Environmental Prediction (NCEP) and National Ocean
 212 Service (NOS) Coast Survey and Development Laboratory (CSDL) jointly operate the
 213 model to produce four-times daily 180-hour water level forecasts (NOAA-OPC, 2023; NOAA-
 214 NOS, 2023). The University of Notre Dame Computational Hydraulics Laboratory also
 215 runs a “shadow” model 168-hour once-daily forecast (Wood, 2023) to compare to ongo-
 216 ing development versions. The model has undergone a suite of improvements, for exam-
 217 ple, by including more levee systems, increasing spatial resolution, refining bottom fric-
 218 tion parameterizations, and increasing both temporal and spatial resolution of atmospheric
 219 forcing (Pe’eri, 2023).

220 STOFS 2D Global is driven by ADCIRC, a community finite-element 2D and 3D
 221 hydrodynamics solver (ADCIRC, 2023) used for modeling tides and coastal storm surge
 222 flooding at local and regional scales (Westerink et al., 2008; Bunya et al., 2010; Hope et
 223 al., 2013). STOFS 2D Global applies ADCIRC’s two-dimensional barotropic solver. To
 224 facilitate solutions on the global domain, the model was updated to include a general-
 225 ized cylindrical mapping system for transforming spherical coordinates to a rectilinear
 226 system. Since the spherical coordinate system has a singularity at the poles, and the cylin-
 227 drical mapping system likewise does not permit elements spanning over the poles, a co-
 228 ordinate rotation is added that places both poles overland and likewise rotates the Cori-
 229 olis, surface wind, and internal wave drag terms (Pringle et al., 2021).

230 The operational implementation is not currently forced with hydrology or with the
 231 ocean’s thermohaline circulation, both of which can impact water levels on longer-term
 232 periods (Pringle et al., 2018). To accommodate these slower time-scale fluctuations, the
 233 modeling system computes the mean water level for the five days prior to the start of
 234 any forecast for both the model and the NOS station measurement, and levels the fore-
 235 cast water levels accordingly. Fast time scale fluctuations associated with hydrology (e.g.
 236 a high intensity local rainfall event in a small scale channel) or with the ocean’s ther-
 237 mohaline system (e.g. the changes in the vertical density structure caused by a passing
 238 hurricane) will appear as additional errors in the forecast. Furthermore, the global model
 239 is not presently coupled to wind wave models and therefore is not forced with wave ra-
 240 diation induced stresses associated with wind wave transformation and breaking. The
 241 wind wave model coupling is typically incorporated into regional models and boosts wa-
 242 ter levels along coasts by between 5 cm and up to 0.5 m in limited regions (Dietrich et
 243 al., 2011; Hope et al., 2013; Joyce, Gonzalez-Lopez, et al., 2019; Joyce, Pringle, et al.,
 244 2019). Again, this will appear as a missing physics bias in the model, although typically
 245 this bias will be highly correlated to the specific coastal geometry and bathymetry and
 246 wind intensity and direction. We apply the five-day prior water level adjustment to the
 247 seven day forecast horizon in this study.

248 **2.1.2 Three-Year Hindcast**

249 A three-year STOFS 2D Global hindcast was run for the period of September 2016
 250 - September 2019 with one month spin-up and a 6-second time step. As discussed in Sec-
 251 tion 3, the hindcast was used to train and validate the ML framework. We exercised AD-
 252 CIRC version 55. The internal tide dissipation and boundary layer dissipation param-
 253 eters were adopted from Blakely et al. (2022). The following options were indicated: ICS=-
 254 22 (Mercator projection with pole rotation), IM=511113 (implicit mode), A00=0.8, B00=0.2,
 255 C00=0.0 (time weighting factors), H0=0.1 (minimum water depth), TAU0=0.05 (Gen-
 256 eralized Wave-Continuity Equation, GWCE, weighting factor that weights the relative
 257 contribution of the primitive and wave portions of the GWCE), NTIP=2 (tidal poten-
 258 tial and self attraction / load tide forcings are used), and DT=6 (simulation time step
 259 in seconds). Eight dominant astronomical tidal harmonic constituents (M_2 , N_2 , S_2 , K_2 ,
 260 K_1 , Q_1 , O_1 , P_1) were forced using the tidal potential function as well as self-attraction
 261 and loading. Additionally, atmospheric forcing was sourced from the NCEP Coupled Fore-

cast System Model version 2 (CFSv2), a reanalysis product with 0.25-deg spatial resolution.

The hindcast was run on the Texas Advanced Computing Center’s Frontera supercomputer (Intel Xeon Platinum 8280, clock rate 2.7Ghz, Peak node performance 4.8TF, double precision).

2.2 ML Framework

2.2.1 General TFT Operation

The TFT models exercised herein were based on the formulation given by Lim et al. (2021). We used the Python library Darts (Herzen et al., 2022) to train and evaluate them. The TFT architecture possesses multiple desirable components for spatiotemporal forecasting tasks. TFTs employ variable selection networks to ensure that relevant input is considered at a given time step. Meanwhile, LSTM layers learn short-term temporal dependencies and allow for heterogeneous data sources to be encoded as static and dynamic covariates. Gating mechanisms let the model ignore architecture components where appropriate, while multi-head attention layers provide the model the ability to learn interpretable long-range time dependencies. While the original TFT implementation uses quantile regression to output probabilistic forecasts, we primarily used a loss function in place of likelihood to make deterministic predictions. The hidden layer size, which is shared across both LSTM and self-attention layers, the number of attention heads, the number of LSTM layers, and the learning rate for our model were tuned in a hyperparameteric sweep discussed later in this section.

The target of each TFT was the signed difference between the observed surface water elevation (from NOAA observed station time histories) and the predicted surface water elevation (from the three-year hindcast). Each TFT forecasted (i.e. decoded) this target seven days (168 hours) into the future with 1-hour temporal resolution. It made this forecast in a “single shot” as opposed to autoregressively. We considered a 120-length (120 hour = 5 days) input or encoding region. To facilitate each TFT making this forecast, we supplied user-defined covariates. We supplied them for a period prior to the forecasting horizon (i.e. past covariates), within the forecasting horizon (i.e. future covariates), and also invariantly of time (i.e. static covariates), Table 1. It is noteworthy that three time covariates (hour of the day, day of the week, and month of the year) were included to facilitate the model learning seasonal, diurnal, and semi-diurnal trends. The static covariates, in turn, were provided to suggest spatial coherence. The physics-based covariates, finally, were considered to facilitate predictions at wind-dominant stations. In a post-processing step extrinsic to each TFT, the TFT’s output was added to ADCIRC’s raw prediction to render a “ML-Corrected” ADCIRC prediction.

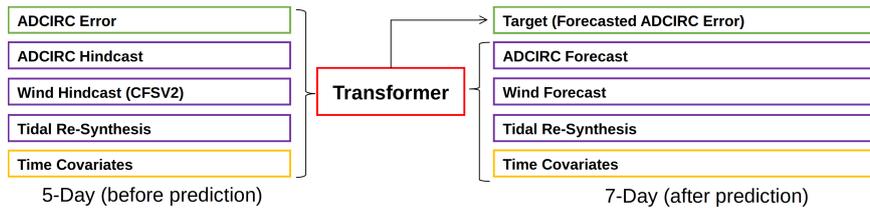
It is noteworthy that the hindcast-centric treatment of the TFT provided herein deviates slightly from how this methodology could be applied operationally for forecasting, Figure 4. Here, we used hindcasted ADCIRC predicted water levels and hindcasted meteorology in place of the forecasted future covariates. Leveraging these hindcasted covariates enabled us to bypass the late-horizon uncertainty stemming from our meteorological forcing. Consequently, the results provided herein indicate intrinsic TFT performance unblemished by spurious meteorological forcing; however, in forecasting mode, this uncertainty would likely lead to degraded late-horizon performance.

2.2.2 Data Processing

The study considered surface water elevations from a three-year ADCIRC-driven STOFS 2D Global hindcast (Section 2.2) and observed surface water elevations at NOAA stations. Both products maintain a 6-minute temporal resolution. The 228 stations, plotted and characterized with regard to tidal or wind dominance in Figure 5, were distributed

Table 1. Past, future, and static covariates considered by the TFT model.

Covariate	Type	Kind	Description
Hour of the Day	Past & Future	Time	Integer from 1-24
Day of the Week	Past & Future	Time	Integer from 1-7
Month of the Year	Past & Future	Time	Integer from 1-12
10-m U Wind Speed (CFSv2)	Past & Future	Dynamic Physics	Float
10-m V Wind Speed (CFSv2)	Past & Future	Dynamic Physics	Float
Tide Resynthesis	Past & Future	Dynamic Physics	Float
η Prediction (ADCIRC)	Past & Future	Dynamic Physics	Float
Latitude	Static	Location	Float
Longitude	Static	Location	Float

**Figure 4.** Diagram of transformer in an operational forecasting scheme.

311 in four regions: (1) Alaska, (2) West: US Pacific Seaboard, Hawaii, Midway Atoll, (3)
 312 Gulf: Gulf of Mexico, and (4) East: US Atlantic Seaboard, Puerto Rico, Bermuda. Most
 313 of the stations are tidally-dominant (i.e. the tidal potential energy was on the order of
 314 the total potential energy at these locations); however, some are situated in strongly wind-
 315 dominated locations (i.e. shallow locations wherein water levels are sensitive to the pre-
 316 vailing winds). The Gulf region is composed almost exclusively of wind-dominant sta-
 317 tions while the East region includes several in the shallow Chesapeake and Delaware Bays.
 318 When we evaluate performance in Section 3, we will discriminate tidally-dominant from
 319 wind-dominant stations.

320 We mapped meteorological forcing (from CFSv2) and model surface water eleva-
 321 tion predictions to these stations during the hindcast simulation. These data were then
 322 extracted at the top of each hour. The corresponding NOAA observed surface water el-
 323 evations at the top of each hour were then matched to these data. Model and observed
 324 water levels were with respect to the mean sea level (MSL) vertical datum. Due to miss-
 325 ing NOAA observed levels for some stations, this matching was incomplete. Given that
 326 we did not fill in missing data, the entirety of the hindcast period for some stations could
 327 not be used for training / testing. Rather, we segregated the matched data into time-
 328 contiguous chunks.

329 We trained each TFT on a series of time-contiguous chunks. The chunking was ac-
 330 complished by first segregating the matched data into time-contiguous spans that were
 331 at least longer than the desired chunk size (i.e. number of hours). For example, the 5-
 332 day encoding region and 7-day decoding region considered herein constituted a chunk
 333 size of 288. Thereafter, for each span, a window with length equal to the chunk size tra-
 334 versed the span incrementally. For each increment, the chunk within the window was ex-
 335 tracted either for training, validation, or testing. It moved chronologically along the span
 336 with a predefined shift length. A shift length equal to the chunk size would result in no
 337 chunk overlap; a shift length less than the chunk size would result in chunk overlap and
 338 more extracted chunks. We adopted a 1-day shift length to guarantee some overlap.

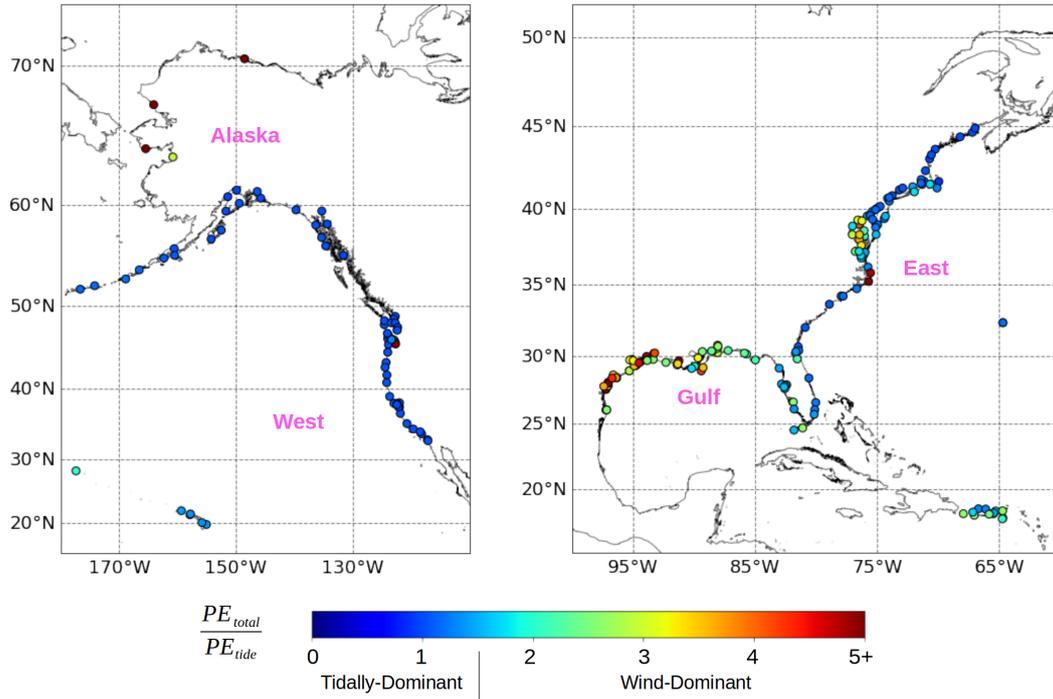


Figure 5. NOAA observed surface water elevation stations considered in study. Stations were grouped into the following regions to facilitate evaluation: (1) Alaska, (2) the US Pacific Seaboard, Hawaii, and Midway Atoll denoted “West”, (3) the Gulf of Mexico denoted “Gulf”, and (4) the US Atlantic Seaboard, Puerto Rico, and Bermuda denoted “East”. Colors reference the ratio between a given station’s total potential energy to its tidal potential energy. Values close to 1 indicate tidal dominance. Higher values suggest wind-dominance.

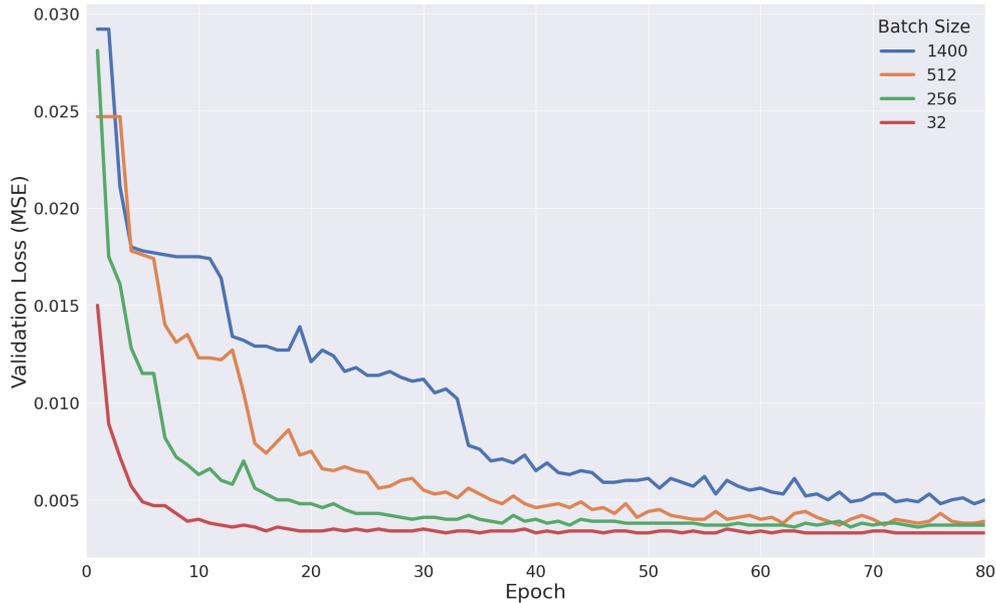


Figure 6. Validation loss plotted as a function of epoch for various batch sizes. Here, validation loss is the average mean squared error over the validation set. Only stations in the Northeast are considered.

339

2.2.3 TFT Training and Hyperparametric Sweep

340

341

342

343

344

345

346

Unless otherwise noted, for each station, we selected the first 70% of the chunks in the hindcast period for training. The remaining chunks were randomly shuffled into validation and test sets (20% validation, 10% testing). This ensured that all of the stations had representation in both training, validation, and testing; however, depending on NOAA data coverage, some stations had more training / validation / testing chunks than others. We trained each TFT on all of the designated training chunks station-by-station in chronological order.

347

348

349

350

351

352

353

354

355

356

357

358

359

360

361

During preliminary studies, we found TFT performance to be largely invariant to batch size (i.e. number of chunks passed to GPU for forward and back propagation at a time), but smaller batch sizes resulted in slightly enhanced performance over a reduced number of epochs (i.e. number of full passes through the training set), Figure 6. We hypothesize that the greater number of gradient updates from smaller batch sizes facilitated optimization. However, smaller batches on performant GPUs required longer training duration. Consequently, we settled on intermediate batch sizes as they yielded similar overall results with reduced wall-clock time. To facilitate the optimizer’s back propagation, we ensured that each batch was composed primarily of chunks from the same region (Figure 5). For example, if the majority of the chunks in a batch were from tidally-dominant stations, the few from less-predictable wind-dominant stations would likely not contribute to that batch’s loss in any significant manner. Batching with same-region chunks was conducted to mitigate this issue. We also interrogated several optimizers including stochastic gradient descent, Adadelta, Adagrad, and quasi-Newton methods, but Adam proved to be the most stable and performant.

362

363

364

365

We adopted the mean squared error (MSE) loss function in place of mean absolute error (MAE) loss to aggressively minimize outliers. In exploratory studies, we observed that both MAE and MSE loss targeted intermediate errors equally effectively, but that MSE loss reduced the larger errors at some wind-dominant stations (e.g. Annapo-

366 lis, Baltimore) more reliably. We could have trained the TFTs with quantile regression
 367 to produce probabilistic forecasts; however, given that our focus herein was rendering
 368 deterministic predictions, we did not pursue this option.

369 Upon settling on a batching strategy, optimizer, and loss function, automated Bayesian
 370 optimization was employed via a Tree-Structured Parzen Estimator (TPE) to efficiently
 371 minimize MSE loss over the following hyperparameters: number of LSTM layers, num-
 372 ber of attention heads, hidden layer dimension, and dropout rate. Our model proved rel-
 373 atively insensitive to these hyperparameters over a subset of stations in the Northeast;
 374 however, modest performance gains were obtained nonetheless. Based on this exercise,
 375 three LSTM layers, three attention heads, and a dropout rate of approximately 10% were
 376 adopted. Hidden layer size was ultimately capped at 110 to keep training time and com-
 377 putational costs manageable.

378 After conducting the hyperparametric sweep, ten models were identified for com-
 379 prehensive training and evaluation. They are described below:

- 380 1. Global - All - Baseline: This model inherited the tuned hyperparameters in ad-
 381 dition to all of the covariates listed in Table 1 (U, V, ADCIRC η , tidal resynthe-
 382 sis and time and spatial covariates). It was trained on the first 70% of each sta-
 383 tion’s chunks in the hindcast period. It was trained on all stations in the Alaska,
 384 West, Gulf, and East regions. In the model, the encoder attends to both past and
 385 future covariates, while the decoder attends only to future covariates.
- 386 2. Global - No Physics: This model was identical to the Baseline model incorporat-
 387 ing the time and spatial covariates, but did not retain any of the physics-based
 388 dynamic covariates (U, V, ADCIRC η , tidal resynthesis).
- 389 3. Global - No Tides: This model was identical to the Baseline model, but was not
 390 given access to the tidal covariates. Using this model, we wanted to assess if tidal-
 391 centric errors could be addressed by the transformer in the absence of a tidal resyn-
 392 thesis.
- 393 4. Global - No Winds: This model was identical to the Baseline model, but lacked
 394 past and future wind covariates.
- 395 5. Global - All - Full Attention: This model derived from the Baseline model, but
 396 adopted a slightly different architecture wherein the decoder was allowed to at-
 397 tend to previous, current, and forthcoming future covariates in the forecast hori-
 398 zon. Enabling the decoder to attend to current and forthcoming future wind data,
 399 and not simply previous covariates as is done in the baseline model, was hypoth-
 400 esized to increase performance in wind-dominant stations.
- 401 6. Global - All - 30% Train: This model derived from the Baseline model, but was
 402 trained on only the first 30% of each station’s chunks in the hindcast period (as
 403 opposed to 70%).
- 404 7. Alaska - All: This model derived from the Baseline model, but was trained exclu-
 405 sively on chunks in the Alaska region.
- 406 8. West - All: This model derived from the Baseline model, but was trained exclu-
 407 sively on chunks in the West region.
- 408 9. Gulf - All: This model derived from the Baseline model, but was trained exclu-
 409 sively on chunks in the Gulf region.
- 410 10. East - All: This model derived from the Baseline model, but was trained exclu-
 411 sively on chunks in the East region.

412 We leveraged Nvidia A6000 GPUs for training and inference. Training a “global” model
 413 required approximately 24-hours of wall-clock time. Evaluating a “global” model’s en-
 414 tire test set for all chunks therein required only a minute of wall-clock time.

3 Results

3.1 Evaluation Metrics

To quantify the performance of each TFT model, we considered four evaluation metrics: normalized root-mean-squared error (NRMSE), maximum error (MAX), coefficient of determination (R^2), and the Willmott skill score (WSS) (Willmott, 1981). They are given by:

$$NRMSE = \frac{\sqrt{\frac{1}{n} \sum_{i=1}^n (\eta_{obs,i} - \eta_{pred,i})^2}}{\sqrt{\frac{1}{n} \sum_{i=1}^n (\eta_{obs,i})^2}} \quad (1)$$

$$MAX = \max(|\eta_{obs,i} - \eta_{pred,i}|) \quad (2)$$

$$R^2 = 1 - \frac{\sum_{i=1}^n (\eta_{obs,i} - \eta_{pred,i})^2}{\sum_{i=1}^n (\eta_{obs,i} - \eta_{obs,avg})^2} \quad (3)$$

$$WSS = 1 - \frac{\sum_{i=1}^n (\eta_{obs,i} - \eta_{pred,i})^2}{\sum_{i=1}^n (|\eta_{pred,i} - \eta_{obs,avg}| + |\eta_{obs,i} - \eta_{obs,avg}|)^2} \quad (4)$$

Note that η_{obs} and η_{pred} denote observed and predicted surface water elevations, respectively, and n denotes output chunk length (168). We consider two sets of predicted water levels: ADCIRC-predicted (i.e. raw ADCIRC) and ML-Corrected (i.e. ADCIRC+ML). NRMSE and MAX are error metrics. Smaller values denote superior model performance. R^2 and WSS are regressive score metrics. For both, a value of 1 denotes perfect skill. These four metrics were used to evaluate individual chunks.

3.2 Evaluated Performance

In this section, we evaluate TFT region-based and station-based performance using the test set of each station. This set was separate from training and validation. Moreover, for each station, the test set period did not overlap with that of the training set. On average, each station had 100 test chunks. Each evaluation considered evaluation metrics calculated for individual chunks, and then these metrics were averaged over regions (see region-based performance) or individual stations (see station-based performance). Figure 7 illustrates how ADIRC + ML improves the solution as compared to ADCIRC alone for one 7-day forecasting chunk at three sample stations.

3.2.1 Region-Based Performance

We calculated the four evaluation metrics for the 7-day forecast horizon in each of the chunks in our test set. Thereafter, we averaged these metrics over all of the chunks at each station. We then grouped these averaged metrics by region. The resulting distributions are plotted in Figure 8. It is clear that all of the TFTs rendered relatively aggressive corrections in Alaska, the West (predominantly Pacific Seaboard), and the East (predominantly Atlantic Seaboard). This is likely because each of these regions have a majority of tidally-dominant stations wherein surface water elevations are determined by highly deterministic, cyclical tidal potential functions and not by lower fidelity, less certain meteorology. The performance of the TFT devoid of physics-based dynamic covariates supports this hypothesis as it rendered relatively aggressive corrections correlating errors predominantly to time covariates.

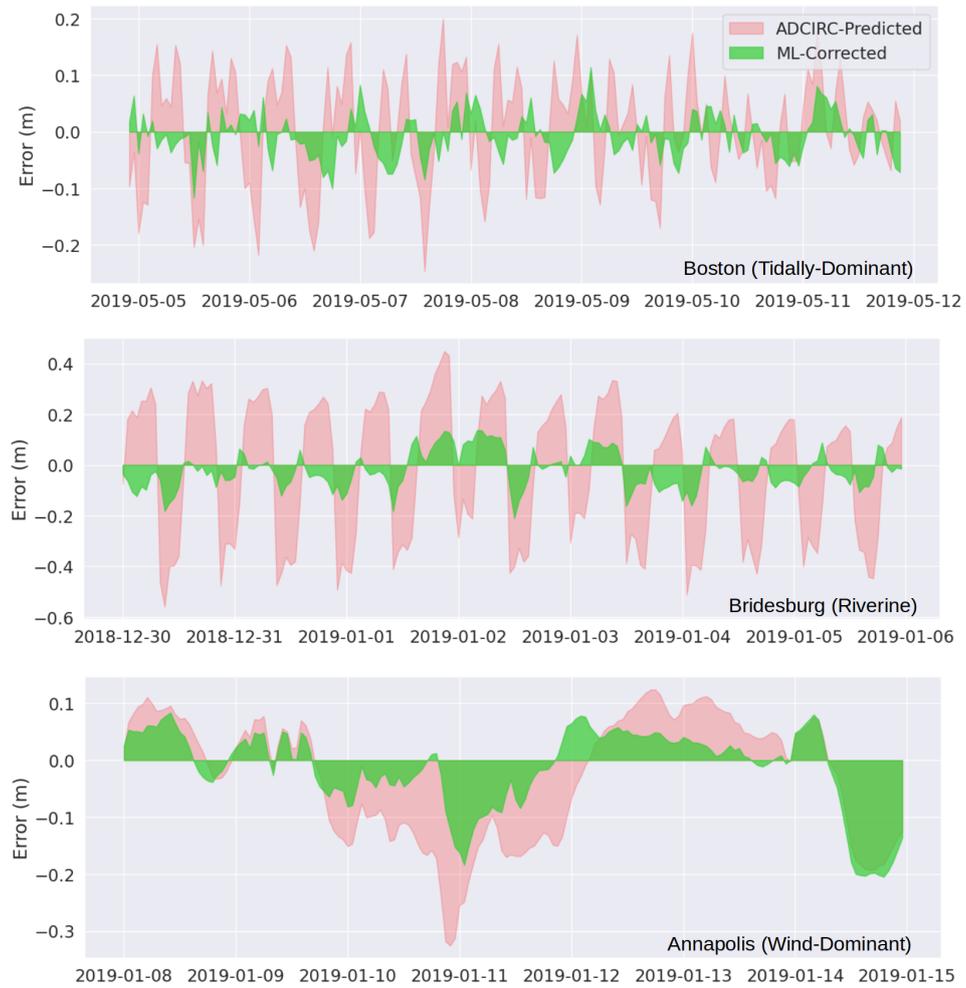


Figure 7. Examples of 7-day output test chunks considered in the evaluation. ADCIRC-predicted and ML-corrected error are considered here. Values closer to x-axis represent higher skill. ML output was produced in a “single shot” for each chunk.

448 In the Gulf of Mexico, however, the TFT-based corrections were less aggressive.
 449 The Gulf has a majority of wind-dominant stations which are difficult to predict, lead-
 450 ing to larger ADCIRC prediction uncertainty and weaker ML-based corrections. Adding
 451 tides and ADCIRC’s predictions (“Global - No Winds”) and thereafter adding wind (“Global
 452 - All - Baseline”) clearly made for more performant TFTs, especially in the Gulf. The
 453 TFT trained with full attention demonstrated roughly the same skill as its counterpart
 454 with limited attention, suggesting that allowing the decoder to attend to current and forth-
 455 coming future covariates does not enhance performance and that the encoder’s atten-
 456 tion to past and future covariates controls model skill. Excluding resynthesized tides but
 457 retaining all other covariates (“Global - No Tides”) rendered similar performance to the
 458 other TFTs supplied with physics-based covariates, suggesting that time covariates are
 459 facilitating tidal pattern recognition in ADCIRC’s error space and surface water eleva-
 460 tion prediction. The TFT with a reduced training set size exhibited approximately the
 461 same skill as the TFT devoid of physics-based dynamic covariates. Finally, the regional
 462 models demonstrated slightly worse performance compared to the “Baseline” model.

463 From this high-level region-centric evaluation, we can assume that (1) the TFT is
 464 intrinsically capable of correcting tidally dominant stations without physics-based dy-
 465 namic covariates, but by adding additional physics-based covariates (viz. ADCIRC pre-
 466 diction), the TFT renders a more aggressive correction and (2) the physics-based covari-
 467 ates facilitate improved TFT performance at wind-dominated stations. We will confirm
 468 these assumptions in the following section wherein we conduct a station-based evalua-
 469 tion.

470 **3.2.2 Station-Based Performance**

471 The evaluation metrics were calculated for each station’s chunks in our test set. First,
 472 we considered TFT late-horizon performance. All time-series forecasting models suffer
 473 from degradation, and this is typically correlated to horizon length. The TFT is no ex-
 474 ception. As shown in Figure 9, TFT skill within the 6-7 day horizon was lower than its
 475 skill within the 0-1 day horizon. Moreover, more prominent drops in late-horizon skill
 476 generally occurred at wind-dominant stations. The “Global - No Physics” TFT exhib-
 477 ited the largest drops while the “Global - All - Baseline” model, with its dynamic wind
 478 covariates, exhibited the lowest. It is noteworthy that even at the most recalcitrant wind-
 479 dominant stations, the degradation of “Global - All - Baseline” was no more than 25%
 480 of the station chunk-averaged NRMSE. In other words, late-horizon degradation, while
 481 quantifiable, was relatively small. It is noteworthy that in a true forecast mode, in the
 482 absence of hindcasted meteorology, the TFT is expected to exhibit more significant late-
 483 horizon degradation since meteorological data for the future seven days will incorporate
 484 forecast uncertainty.

485 To summarize these results, in Figure 10 we plot the range of the first-day and sixth-
 486 day NRMSE averaged over all available test chunks at each station for both ADCIRC
 487 and the “Global - All - Baseline” TFT. Referring to Alaska, it is clear that the TFT was
 488 able to render aggressive corrections with minimal long-horizon degradation at the tidally-
 489 dominant stations. At the region’s four wind-dominant stations, however, the corrections
 490 were comparatively weaker and exhibited more significant long-horizon degradation. This
 491 same trend was observed in the West and East regions. In the East region, some hybrid
 492 stations (tidally-dominant stations with water levels occasionally influenced by winds)
 493 exhibited weaker corrections and moderate late-horizon degradation. In the Gulf, while
 494 no station exhibited significant late-horizon degradation, no aggressive corrections were
 495 rendered save for at the Coast Guard Station in Mobile, Alabama (the TFT dropped the
 496 NRMSE from 0.67 to 0.32).

497 We then considered chunk-averaged performance of the “Global - All - Baseline”
 498 TFT at all 228 stations. Here, we grouped chunks by station and then averaged. Refer-

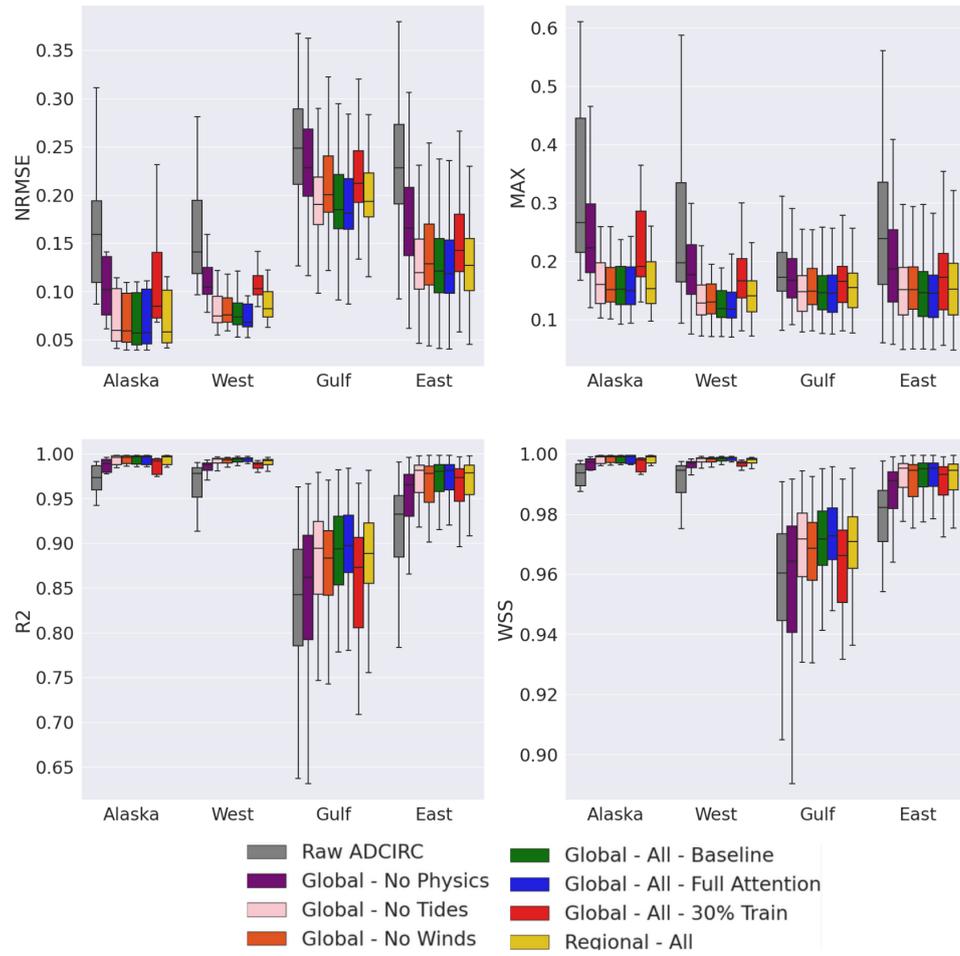


Figure 8. Boxplots of station chunk-averaged evaluation metrics separated by region for ADCIRC and various TFTs. The ML-corrected ADCIRC prediction was considered for each TFT.

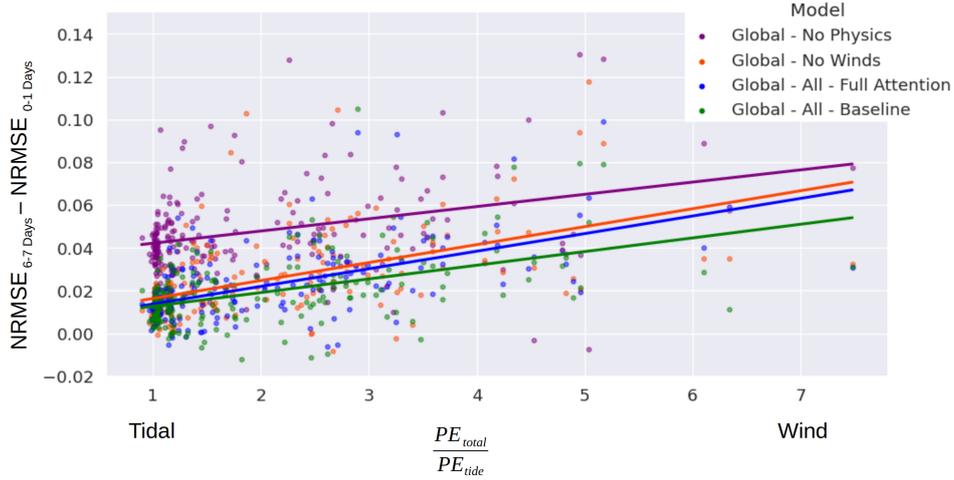


Figure 9. Relationship between late-horizon degradation in ML performance as quantified by $NRMSE_{6-7\ days} - NRMSE_{0-1\ days}$ and tidal-wind dominance as quantified by the ratio between total potential energy and tidal potential energy. In general, every ML model considered indicated late-horizon degradation for both tidally-dominant and wind-dominant stations; however, this degradation was more pronounced at wind-dominant stations. The lines are best-fits to station chunk-averaged degradation for each ML model. All stations were considered.

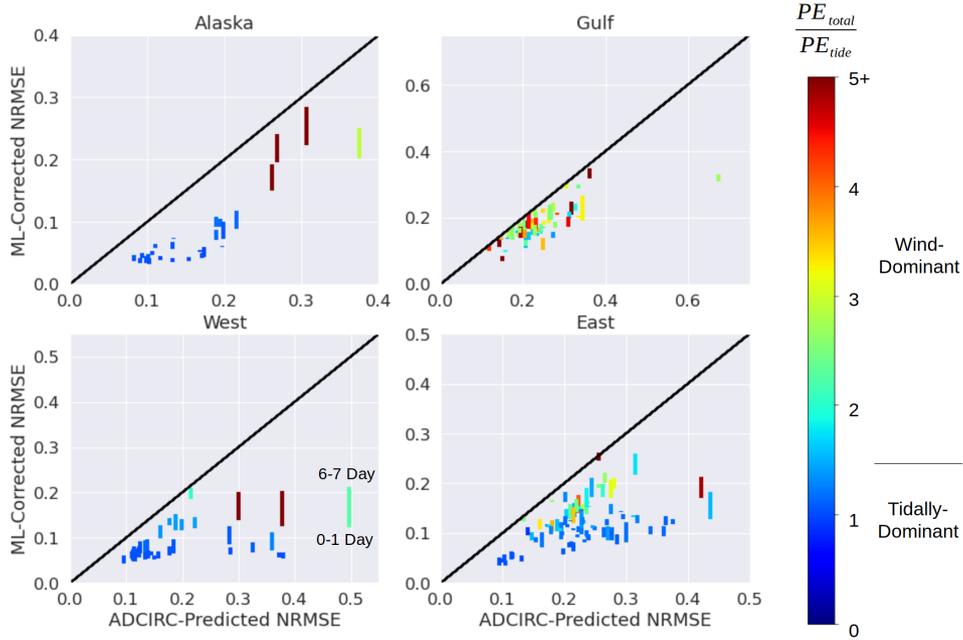


Figure 10. 45-deg plots of the range of ML-corrected vs. ADCIRC-predicted NRMSE for each station. The Global - All - Baseline TFT was used to generate each plot. Each bar corresponds to a station (test chunk-averaged data). The length of each bar indicates the degradation in model skill over the horizon (bottom corresponds to 0-1 day horizon NRMSE, top corresponds to 6-7 day horizon NRMSE). In general, tidally-dominant stations were associated with aggressive corrections and low degradation in skill over the horizon. Wind-dominant stations were associated with comparatively weaker corrections and higher degradation in skill over the horizon.

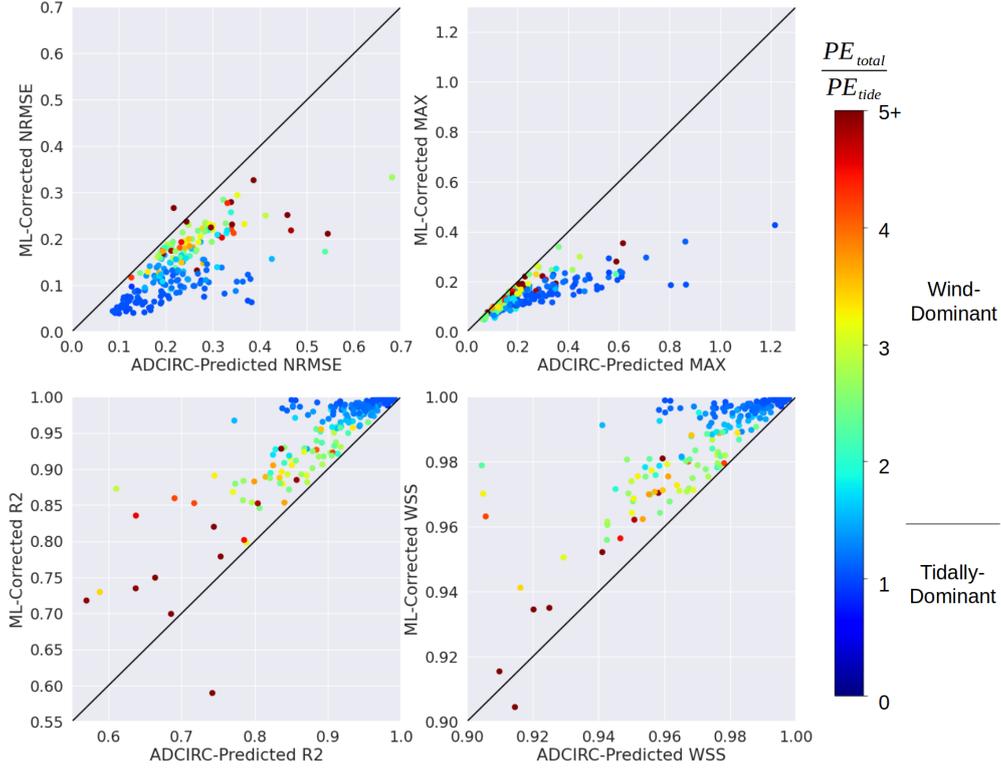


Figure 11. 45-deg plots of ML-corrected vs. ADCIRC-predicted evaluation metrics at the 228 stations throughout the test period. The evaluated ML model was Global - All - Baseline. Each marker corresponds to a station with the metric averaged over all chunks for that station. Colors reference the ratio between a given station’s total potential energy to its tidal potential energy, an indication of tidal to wind dominance.

ring to Figure 11, the ML-corrected chunk-averaged station-specific evaluation metrics vs. the ADCIRC-predicted chunk-averaged station-specific evaluation metrics, it is apparent that the TFT was able to render corrections at all but one station (Copano Bay, Texas). The most aggressive corrections occurred at the tidally-dominant stations; 90% of tidal stations saw NRMSE decrease by at least 25% while 50% saw their NRMSE more than halved. The TFT was able to render corrections at wind-dominant stations; however, they were unable to produce the skill observed at the tidally-dominant stations. Regardless, 40% of wind-dominant stations saw their NRMSE decrease by at least 25% while 65% of stations whose ADCIRC NRMSE was greater than 0.4 also saw their NRMSE more than halved.

To clarify this trend further, we investigated six stations in detail: two tidally-dominant (Boston and Anchorage), two wind-dominant (Annapolis, Baltimore), and two riverine (Bridesburg, Pilottown). Figure12 shows 45-deg plots of predicted vs. observed hourly η values for all test chunks for standalone ADCIRC and four TFT-enhanced forecasts with increasing levels of sophistication: the TFT without any physics-based dynamic covariates (“Global - No Physics”), the TFT with all physics-based dynamic covariates except for winds (“Global - No Winds”), the TFT with all covariates (“Global - All - Baseline”), and the region-centric TFTs with all covariates. It is clear that of these four TFTs, “Global - No Physics” was the least performant. It did make modest corrections at the tidally-dominant and riverine stations, suggesting that it picked up on cyclic patterns in the error space via the time covariates, but it was largely unrespon-

520 sive at wind-dominant stations. It is noteworthy that this model corrected a phasing issue
 521 at Bridesburg and partially corrected the anisotropic bias at Pilottown. Broadly, the
 522 TFTs appear to be capable of correcting over-damping and under-damping behavior and
 523 even phase lags without physics covariates. The “Global - No Winds” TFT was gener-
 524 ally more performant with the addition of tidal and ADCIRC covariates. Most notably,
 525 it corrected entirely the skewed bias produced by ADCIRC at Pilottown. Adding winds
 526 slightly deteriorated performance at the tidally-dominant stations, but significantly im-
 527 proved TFT skill at Annapolis, Baltimore, and Bridesburg. At the tidally-dominant sta-
 528 tions, this behavior was likely caused by the TFT attending to winds that were other-
 529 wise inconsequential. The inclusion of the tidal covariate was meant to help the TFTs
 530 discern the non-tidal contribution of surface water elevations at the stations, but this
 531 inclusion did not improve solutions at either tidally-dominant and wind-dominant sta-
 532 tions. Including winds in the “Global - All - Baseline” TFT rendered improved perfor-
 533 mance at Annapolis and Baltimore; however, at the riverine Pilottown station, the ad-
 534 dition of winds actually deteriorated performance. Except at the Annapolis station, the
 535 region-based models were generally less performant than their global counterpart. This
 536 suggests that the TFTs clearly benefited from training on a multitude of signals regard-
 537 less of region.

538 To explore the entitlement of adding wind covariates, we plot test-chunk-based NRM-
 539 SEs for the six stations considered above in Figure 13. Here, we exercised only “Global
 540 - All - Baseline” and “Global - No Winds”. For the tidally-dominant stations (Boston
 541 and Anchorage), it is apparent that both transformers produced predictions of similar
 542 skill. Annapolis and Baltimore, both wind-dominant stations, saw marked improvements
 543 from adding wind covariates. Moreover, it is clear that large ADCIRC errors were gener-
 544 ally associated with high-wind events, suggesting that attending to winds results in
 545 enhanced transformer correction capacity. The riverine station Pilottown saw a marginal
 546 improvement in chunk-based NRMSE with the inclusion of wind covariates. Finally, the
 547 other riverine station, Bridesburg, exhibited an aggressive correction from both TFTs
 548 and demonstrated, for a few high-speed wind chunks, the utility of adding wind covari-
 549 ates.

550 Based on these results, it is apparent that the addition of physics-based dynamic
 551 covariates yields enhanced performance at wind-dominant stations. The inclusion of time
 552 covariates alone produced considerable improvement at tidally-dominant stations, but
 553 failed to render desired performance at the wind-dominant stations. Adding tides did
 554 not enhance transformer performance; however, adding ADCIRC’s own prediction (“Global
 555 - No Winds”) further enhanced skill at both station types, but it was only after adding
 556 wind covariates that more consistent corrections were rendered at wind-dominant sta-
 557 tions (e.g. Annapolis, Baltimore).

558 4 Conclusions and Discussion

559 We have coupled the temporal fusion transformer to a high-fidelity global ocean
 560 hydrodynamics model, STOFS 2D Global, to render improved station-based predictions
 561 of surface water elevation seven days into the future. The STOFS 2D Global model by
 562 itself has generally exhibited high skill along US coastlines and adjacent inland coastal
 563 waters, but nevertheless, it has systemic model discrepancy stemming from inadequate
 564 geometric representation, coarse mesh resolution, incorrect bathymetry, uncertainty in
 565 the parameterization of dissipative processes, and meteorological error. Weakly coupling
 566 a transformer to this physics-based hydrodynamics model enables the identification of
 567 patterns in model error space and their subsequent reduction. We considered several trans-
 568 formers in this study, each supplied with different covariates to ascertain *what* param-
 569 eters lead to a skillful corrector. We trained and evaluated each transformer on a three-
 570 year ADCIRC STOFS 2D Global hindcast. The transformers produced 7-day predictions
 571 of ADCIRC error in a “single shot” with 1-hour temporal resolution, ingesting five days

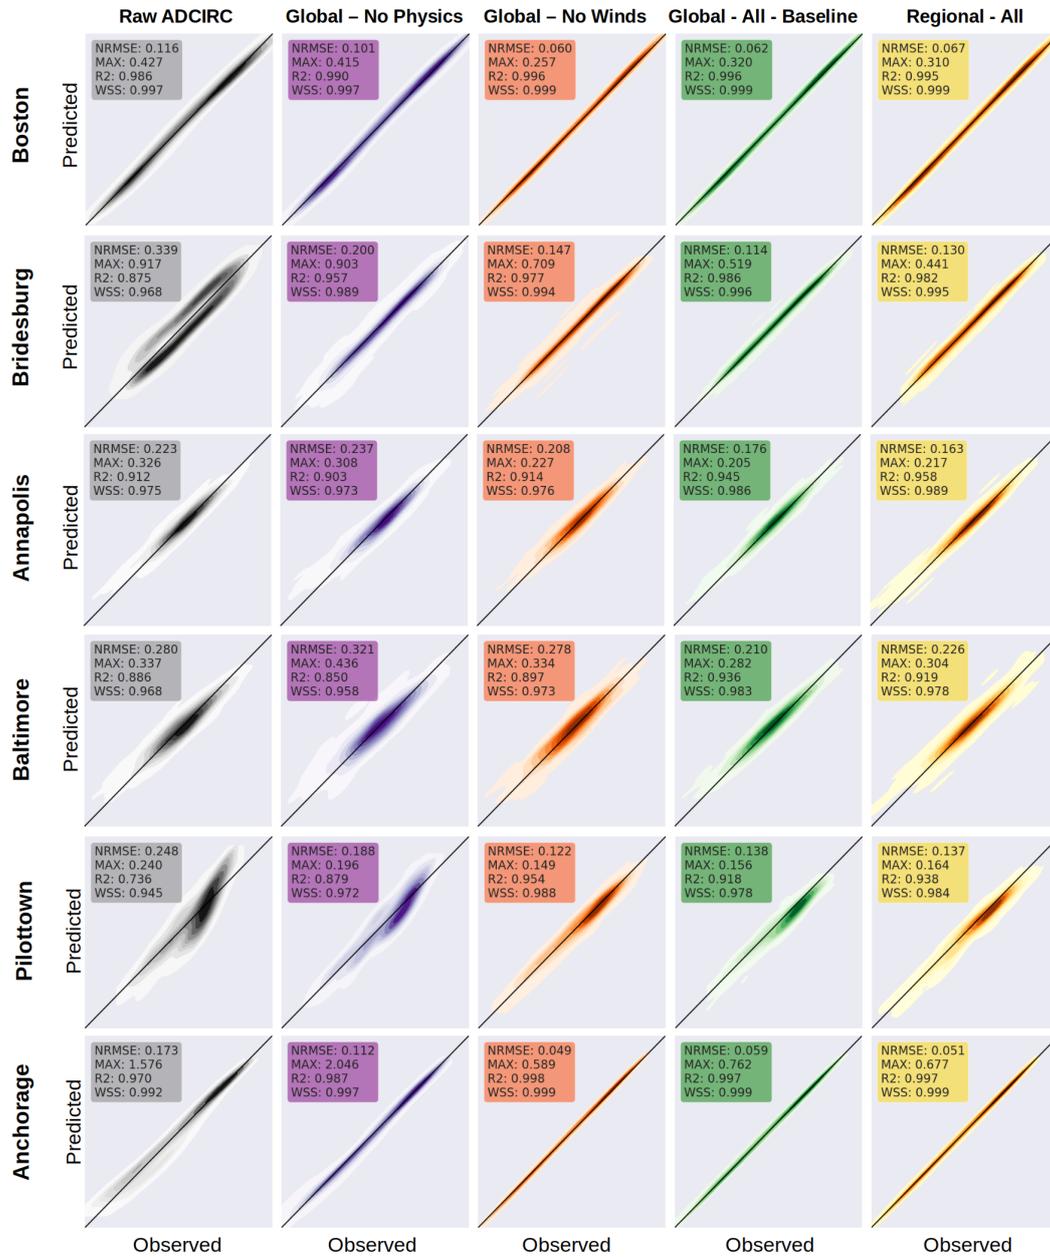


Figure 12. 45-deg plots of predicted vs. observed surface water elevation for hourly data for all test forecast periods. Values from the 6-7 day forecast horizon are used. Annapolis and Baltimore are wind-dominant stations in the Chesapeake Bay (East), Bridesburg is a moderate wind-dominant station located in the Delaware River (East), Anchorage is a tidally dominant station in Alaska, and Pilottown is a wind-dominant station located in Louisiana (Gulf). The ML-corrected ADCIRC prediction was considered for each TFT.

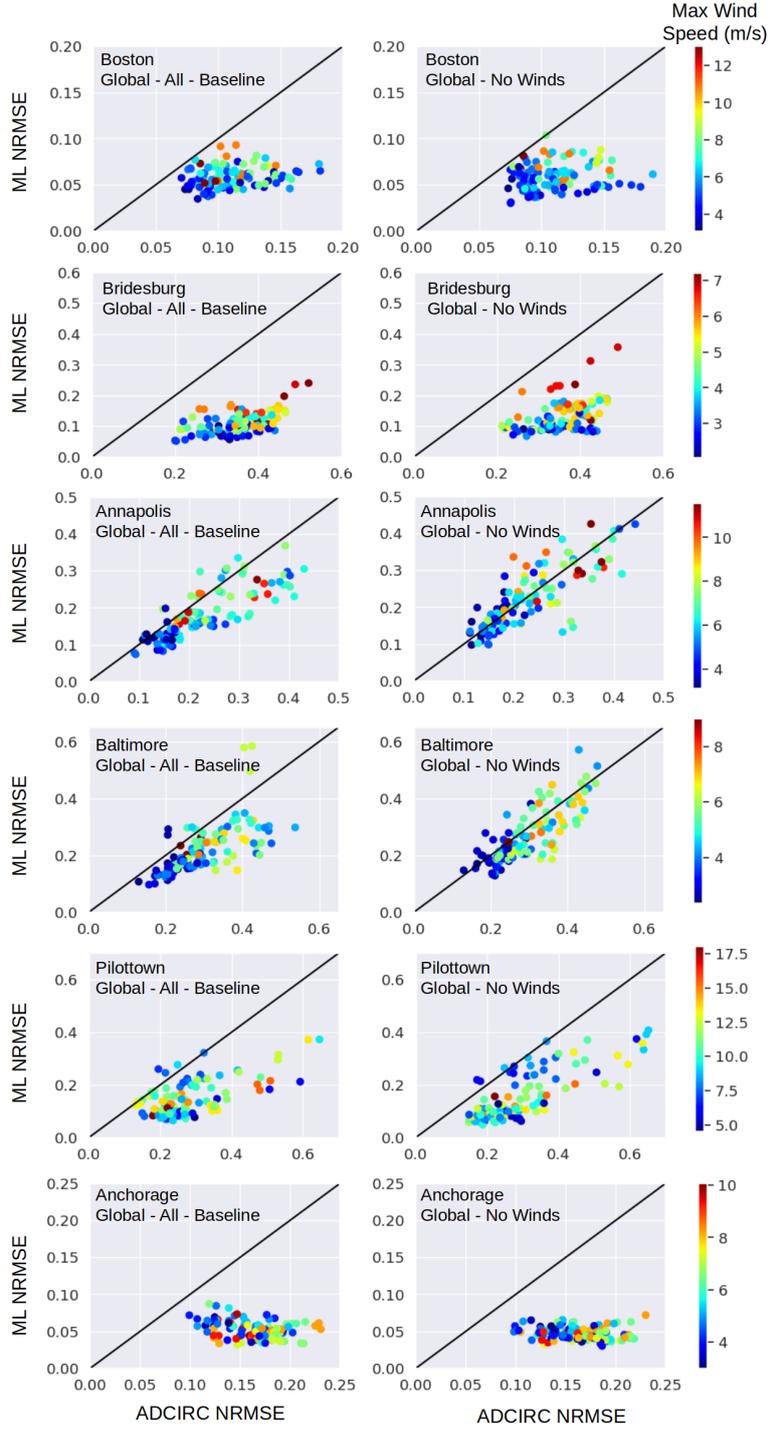


Figure 13. 45-deg plots of ML-corrected vs. ADCIRC-predicted NRMSE averaged for each test chunk at the station. The left column considers Global - All - Baseline, which was trained on winds. The right column considers Global - No Winds, which was not trained on winds. Each marker corresponds to a 7-day output test chunk. In general, the wind-trained transformer was more performant at wind-dominant stations than the transformer trained without winds. Performance at tidally-dominant stations was largely insensitive to winds.

572 of data prior to the forecast horizon. Based on the evaluation of each transformer, we
 573 conclude the following:

- 574 • In general, a single transformer exhibits sufficient skill to consistently correct sur-
 575 face water elevations at hundreds of stations along the US coastline.
- 576 • Each transformer was particularly capable at tidally-dominant stations. Even in
 577 the absence of physics-based dynamic covariates, the TFT was able to render ag-
 578 gressive corrections. For the best-performing transformer, which incorporated AD-
 579 CIRC computed water levels and wind covariates throughout the 5-day hindcast
 580 and 7-day forecast, 50% of tidal stations saw their NRMSE (averaged over the test
 581 period) halved. In certain cases, the inclusion of wind covariates slightly deteri-
 582 orated performance at tidally-dominant stations. This suggests that the TFT was
 583 attending to winds, even in cases when ADCIRC error was wind-invariant.
- 584 • At wind-dominant stations, adding more physics-based dynamic covariates led to
 585 enhanced skill. Including none of these resulted in little to no corrections. Adding
 586 ADCIRC’s predictions certainly improved performance at wind-dominant stations;
 587 however, wind covariates were necessary to correct recalcitrant test chunks, espe-
 588 cially those with high wind-speed events. With the full complement of physics-
 589 based covariates, 40% of wind-dominant stations saw their NRMSE reduced by
 590 at least 25%.
- 591 • Transformers were either trained on all stations along US coastlines or on stations
 592 in specific regions. Region-centric transformers exhibited slightly diminished per-
 593 formance compared to their counterparts trained on all available stations. The wind-
 594 dominant Annapolis station was one exceptional station wherein the region-centric
 595 model out-performed the global transformers.
- 596 • Of all the transformers considered, a TFT trained on 2-years of the hindcast with
 597 the full complement of past, future, and static covariates was, with a few excep-
 598 tions, the most performant. This TFT was denoted “Global - All - Baseline” herein.
 599 Exceptions include the aforementioned slight deterioration at tidally-dominant sta-
 600 tions (owing to needlessly attending to winds) and rare instances of region-based
 601 TFT superiority. The TFT trained with full attention, so that its decoder could
 602 attend to forthcoming future covariates, exhibited roughly similar performance with
 603 a larger training cost as compared to the same model whose decoder could only
 604 attend to previous future covariates. Both of these models had encoders that did
 605 attend to both past and future covariates.
- 606 • The TFTs exhibited modest amounts of late-horizon degradation in skill. This degra-
 607 dation was measured, test chunk by test chunk, as the difference between the NRMSE
 608 for the 6-7 day horizon and the NRMSE for the 0-1 day horizon. This degrada-
 609 tion was noted to be the highest at some wind-dominant stations; however, it never
 610 exceeded 25% of a given station’s average NRMSE.

611 The approach proposed herein is station-centric. Training and evaluation data was
 612 mapped to hundreds of stations situated along US coastlines, and improvements were
 613 rendered by the transformers at these specific locations. The extrapolative capability of
 614 the TFT was not assessed. While it could theoretically be used to extrapolate beyond
 615 the trained stations, it is likely that the training set size would need to include thousands
 616 (and not hundreds) of stations so that the TFT could draw correlations between the sup-
 617 plied location-based static covariates, the physics-based dynamic covariates, and the tar-
 618 get signal.

619 The framework was challenged to attend to high wind-speed events. In the present
 620 work, the vast majority of chunks did not incorporate elevated wind levels and their as-
 621 sociated elevated surges. The few chunks that did were not accommodated in any par-
 622 ticular fashion. In fact, while MSE loss penalized outliers aggressively, it did so in batches.
 623 Consequently, the few chunks with fringe events were likely muted by the more frequent

624 quiescent chunks. There are potential pathways to circumvent this undesirable behav-
 625 ior. For example, in imbalanced classification problems, class weights make the optimizer
 626 more cognizant of under-represented classes. By tagging our fringe chunks and ampli-
 627 fying their loss, we hypothesize that the optimizer may become more sensitive to them.
 628 However, this does not address the shortcomings of the wind product. CFSv2 has 0.25
 629 deg spatial resolution. Consequently, it is not expected to perform well in regions with
 630 variable terrain and complex shallow inland water systems whose water levels are par-
 631 ticularly susceptible to strong winds. Additionally, this lack of resolution means that it
 632 cannot adequately resolve strong frontal systems and high-energy, low-pressure wind events
 633 such as tropical cyclones which tend to be muted in the product. Finally, it does not have
 634 skillful inland atmospheric boundary layer adjustments, further reducing its utility for
 635 the inland and near-shore stations considered herein.

636 Finally, maturing the proposed framework for operational forecasting will require
 637 adopting forecasted meteorology and water level predictions in place of hindcasted val-
 638 ues. This will introduce a greater level of epistemic uncertainty. Herein, we trained our
 639 TFTs using a deterministic loss function. This deterministic model could be used to fa-
 640 cilitate forward propagation in an ensemble uncertainty quantification scheme. For ex-
 641 ample, each member of NOAA’s Global Ensemble Forecast System (GEFS) could be forced
 642 independently through the transformer. Thereafter, probabilities of exceedance could be
 643 calculated at each station. Additionally, leveraging quantile regression in place of a de-
 644 terministic loss function could help quantify confidence of each correction in time. Col-
 645 lectively, this strategy offers a compelling pathway to operations which are becomingly
 646 increasingly stochastic in nature.

647 Open Research Section

648 Results from the three-year ADCIRC STOfS 2D Global hindcast used to train the
 649 transformers considered herein can be downloaded from the “Improving Storm Surge Fore-
 650 casts with Transformers” project (Cerrone et al., 2023) on DesignSafe (NSF-NHERI, 2023).
 651 Moreover, output from each transformer can be found under the same DesignSafe project.
 652 This project can be accessed here: <https://doi.org/10.17603/ds2-t5mf-3757>

653 Acknowledgments

654 The authors would like to acknowledge Drs. Rick Luettich and Shintaro Bunya, both
 655 of the University of North Carolina at Chapel Hill, and Dr. Eirik Valseth, of the Nor-
 656 wegian University of Life Sciences and the Oden Institute for Computational Engineer-
 657 ing and Sciences at the University of Texas at Austin, for their valuable feedback dur-
 658 ing the early stages of this work. This work was funded by the United States Depart-
 659 ment of Energy (DOE) award DE-SC0022316. The three-year hindcast was run on TACC’s
 660 Frontera supercomputer with allocation “High Fidelity Hurricane Storm Surge and Ocean
 661 Modeling” (DMS21031).

662 References

- 663 *Adcirc.* (2023). Retrieved 22.11.2023, from `\url{https://adcirc.org/}`
- 664 Ayyad, M., Hajj, M. R., & Marsooli, R. (2022, November). Machine learning-based
 665 assessment of storm surge in the new york metropolitan area. *Sci. Rep.*, *12*(1),
 666 1–12.
- 667 Beaman, R. (2018). High-resolution depth model for the northern australia—100 m.
 668 *Geoscience Australia, Canberra.*
- 669 Blakely, C. P., Ling, G., Pringle, W. J., Contreras, M. T., Wirasaet, D., Westerink,
 670 J. J., . . . Massey, C. (2022, May). Dissipation and bathymetric sensitivities
 671 in an unstructured mesh global tidal model. *J. Geophys. Res. C: Oceans,*

- 672 127(5).
- 673 Bolton, T., & Zanna, L. (2019, January). Applications of deep learning to ocean
674 data inference and subgrid parameterization. *J. Adv. Model. Earth Syst.*,
675 11(1), 376–399.
- 676 Bonavita, M., & Laloyaux, P. (2020, December). Machine learning for model er-
677 ror inference and correction. *Journal of Advances in Modeling Earth Systems*,
678 12(12), e2020MS002232.
- 679 Bunya, S., Dietrich, J. C., Westerink, J. J., Ebersole, B. A., Smith, J. M., Atkinson,
680 J. H., . . . Roberts, H. J. (2010, February). A High-Resolution coupled riverine
681 flow, tide, wind, wind wave, and storm surge model for southern louisiana and
682 mississippi. part i: Model development and validation. *Mon. Weather Rev.*,
683 138(2), 345–377.
- 684 Butler, T., Altaf, M. U., Dawson, C., Hoteit, I., Luo, X., & Mayo, T. (2012). Data
685 assimilation within the advanced circulation (adcirc) modeling framework
686 for hurricane storm surge forecasting. *Monthly Weather Review*, 140(7),
687 2215 - 2231. Retrieved from [https://journals.ametsoc.org/view/
688 journals/mwre/140/7/mwr-d-11-00118.1.xml](https://journals.ametsoc.org/view/journals/mwre/140/7/mwr-d-11-00118.1.xml) doi: [https://doi.org/10.1175/
689 MWR-D-11-00118.1](https://doi.org/10.1175/MWR-D-11-00118.1)
- 690 Butler, T., Graham, L., Estep, D., Dawson, C., & Westerink, J. (2015). Definition
691 and solution of a stochastic inverse problem for the manning’s n parameter
692 field in hydrodynamic models. *Advances in Water Resources*, 78, 60-79.
693 Retrieved from [https://www.sciencedirect.com/science/article/pii/
694 S0309170815000135](https://www.sciencedirect.com/science/article/pii/S0309170815000135) doi: <https://doi.org/10.1016/j.advwatres.2015.01.011>
- 695 Cerrone, A., Westerink, L., Dawson, C. N., & Westerink, J. (2023). *Improving
696 storm surge forecasts with transformers*. [Dataset] Designsafe-CI. Re-
697 trieved from [https://www.designsafe-ci.org/data/browser/public/
698 designsafe.storage.published/PRJ-4200](https://www.designsafe-ci.org/data/browser/public/designsafe.storage.published/PRJ-4200) doi: 10.17603/DS2-T5MF-3757
- 699 Chen, C., Beardsley, R. C., Luettich Jr., R. A., Westerink, J. J., Wang, H., Per-
700 rie, W., . . . Toulany, B. (2013). Extratropical storm inundation testbed:
701 Intermodel comparisons in scituate, massachusetts. *Journal of Geophys-
702 ical Research: Oceans*, 118(10), 5054-5073. Retrieved from [https://
703 agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20397](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20397) doi:
704 <https://doi.org/10.1002/jgrc.20397>
- 705 De Kleermaeker, S., Verlaan, M., Mortlock, T., Rego, J., Irazoqui, M., Yan, K., &
706 Twigt, D. (2017, 06). Global-to-local scale storm surge modelling on tropical
707 cyclone affected coasts..
- 708 de Oliveira, M. M. F., Ebecken, N. F. F., de Oliveira, J. L. F., & de Azevedo Santos,
709 I. (2009, January). Neural network model to predict a storm surge. *J. Appl.
710 Meteorol. Climatol.*, 48(1), 143–155.
- 711 Dietrich, J. C., Westerink, J. J., Kennedy, A. B., Smith, J. M., Jensen, R. E., Zi-
712 jlema, M., . . . Cobell, Z. (2011). Hurricane gustav (2008) waves and storm
713 surge: Hindcast, synoptic analysis, and validation in southern louisiana.
714 *Monthly Weather Review*, 139(8), 2488 - 2522. Retrieved from [https://
715 journals.ametsoc.org/view/journals/mwre/139/8/2011mwr3611.1.xml](https://journals.ametsoc.org/view/journals/mwre/139/8/2011mwr3611.1.xml)
716 doi: <https://doi.org/10.1175/2011MWR3611.1>
- 717 Elfring, J., Torta, E., & van de Molengraft, R. (2021, January). Particle filters: A
718 Hands-On tutorial. *Sensors*, 21(2).
- 719 Fisheries, & Canada, O. (2023). *Canadian Hydrographic Service Non-Navigational
720 (NONNA) Bathymetric Data [10m resolution]*. Abacus Data Network. Re-
721 trieved from <https://hdl.handle.net/11272.1/AB2/YJPER2> doi: 11272.1/
722 AB2/YJPER2
- 723 Hager, W. H., Castro-Orgaz, O., & Hutter, K. (2019). Correspondence between
724 de saint-venant and boussinesq. 1: Birth of the shallow–water equations.
725 *Comptes Rendus Mécanique*, 347(9), 632-662. Retrieved from [https://
726 www.sciencedirect.com/science/article/pii/S1631072119301263](https://www.sciencedirect.com/science/article/pii/S1631072119301263) doi:

- 727 <https://doi.org/10.1016/j.crme.2019.08.004>
- 728 Haidvogel, D. B., Arango, H. G., Hedstrom, K., Beckmann, A., Malanotte-Rizzoli,
729 P., & Shchepetkin, A. F. (2000). Model evaluation experiments in the north
730 atlantic basin: simulations in nonlinear terrain-following coordinates. *Dynam-*
731 *ics of Atmospheres and Oceans*, 32(3), 239-281. Retrieved from [https://](https://www.sciencedirect.com/science/article/pii/S037702650000049X)
732 www.sciencedirect.com/science/article/pii/S037702650000049X doi:
733 [https://doi.org/10.1016/S0377-0265\(00\)00049-X](https://doi.org/10.1016/S0377-0265(00)00049-X)
- 734 Hervouet, J.-M. (2007). *Hydrodynamics of free surface flows: Modelling with the fi-*
735 *nite element method*. Chichester, West Sussex: John Wiley & Sons Ltd.
- 736 Herzen, J., Lässig, F., Piazzetta, S. G., Neuer, T., Tafti, L., Raille, G., ... Grosch,
737 G. (2022). Darts: User-Friendly modern machine learning for time series. *J.*
738 *Mach. Learn. Res.*, 23(124), 1–6.
- 739 Holland, M. (2020). *An introduction to the extended kalman filter*. Nova Science
740 Publishers.
- 741 Hope, M. E., Westerink, J. J., Kennedy, A. B., Kerr, P. C., Dietrich, J. C., Daw-
742 son, C., ... Westerink, L. G. (2013, September). Hindcast and validation of
743 hurricane ike (2008) waves, forerunner, and storm surge. *J. Geophys. Res. C:*
744 *Oceans*, 118(9), 4424–4460.
- 745 IHO-UNESCO. (2020). *Gebco_2020 grid* (No. DOI: 10.5285/a29c5465-b138-234d-
746 e053-6c86abc040b9). ([https://www.gebco.net/data_and_products/gridded](https://www.gebco.net/data_and_products/gridded_bathymetry_data/gebco-2020/)
747 [_bathymetry_data/gebco-2020/](https://www.gebco.net/data_and_products/gridded_bathymetry_data/gebco-2020/))
- 748 Joyce, B. R., Gonzalez-Lopez, J., Van der Westhuysen, A. J., Yang, D., Pringle,
749 W. J., Westerink, J. J., & Cox, A. T. (2019). U.s. iios coastal and
750 ocean modeling testbed: Hurricane-induced winds, waves, and surge for
751 deep ocean, reef-fringed islands in the caribbean. *Journal of Geophys-*
752 *ical Research: Oceans*, 124(4), 2876-2907. Retrieved from [https://](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JC014687)
753 agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JC014687 doi:
754 <https://doi.org/10.1029/2018JC014687>
- 755 Joyce, B. R., Pringle, W. J., Wirasaet, D., Westerink, J. J., Van der Westhuysen,
756 A. J., Grumbine, R., & Feyen, J. (2019). High resolution modeling of west-
757 ern alaskan tides and storm surge under varying sea ice conditions. *Ocean*
758 *Modelling*, 141, 101421. Retrieved from [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1463500318303834)
759 [science/article/pii/S1463500318303834](https://www.sciencedirect.com/science/article/pii/S1463500318303834) doi: [https://doi.org/10.1016/](https://doi.org/10.1016/j.ocemod.2019.101421)
760 [j.ocemod.2019.101421](https://doi.org/10.1016/j.ocemod.2019.101421)
- 761 Kerr, P. C., Donahue, A. S., Westerink, J. J., Luettich Jr., R. A., Zheng, L. Y.,
762 Weisberg, R. H., ... Cox, A. T. (2013). U.s. iios coastal and ocean mod-
763 eling testbed: Inter-model evaluation of tides, waves, and hurricane surge in
764 the gulf of mexico. *Journal of Geophysical Research: Oceans*, 118(10), 5129-
765 5172. Retrieved from [https://agupubs.onlinelibrary.wiley.com/doi/abs/](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20376)
766 [10.1002/jgrc.20376](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20376) doi: <https://doi.org/10.1002/jgrc.20376>
- 767 Kerr, P. C., Martyr, R. C., Donahue, A. S., Hope, M. E., Westerink, J. J., Luet-
768 tich Jr., R. A., ... Westerink, H. J. (2013). U.s. iios coastal and ocean
769 modeling testbed: Evaluation of tide, wave, and hurricane surge response
770 sensitivities to mesh resolution and friction in the gulf of mexico. *Jour-*
771 *nal of Geophysical Research: Oceans*, 118(9), 4633-4661. Retrieved from
772 <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/jgrc.20305>
773 doi: <https://doi.org/10.1002/jgrc.20305>
- 774 Laplace, P. (1776). Recherches sur plusieurs points du système du monde. *Mémoires*
775 *de l'Académie (royale) des sciences de l'Institut (imperial) de France*, 89,
776 177–264.
- 777 Li, C., Mahadevan, S., Ling, Y., Wang, L., & Choze, S. (2017, January). A dynamic
778 bayesian network approach for digital twin. In *19th AIAA Non-Deterministic*
779 *approaches conference*. Reston, Virginia: American Institute of Aeronautics
780 and Astronautics.

- 781 Lim, B., Arik, S. Ö., Loeff, N., & Pfister, T. (2021). *Temporal fusion transformers*
782 *for interpretable multi-horizon time series forecasting* (Vol. 37) (No. 4).
- 783 Lyons, M., Larsen, K., & Skone, M. (2022, June). *CoralMapping/AllenCoralAtlas:*
784 *DOI for paper at v1.3*. Zenodo. Retrieved from [https://doi.org/10.5281/](https://doi.org/10.5281/zenodo.6622015)
785 [zenodo.6622015](https://doi.org/10.5281/zenodo.6622015) doi: 10.5281/zenodo.6622015
- 786 NOAA-EMC. (2023). *The global forecast system*. Retrieved from [https://www.emc](https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)
787 [.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php](https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)
- 788 NOAA-NOS. (2023). *Storm surge & tide operational forecast system*. Retrieved from
789 <https://polar.ncep.noaa.gov/estofs/global/index.htm>
- 790 NOAA-OPC. (2023). *Global stofs atlantic region storm surge model guidance*.
791 Retrieved from [https://ocean.weather.gov/estofs/estofs_surge_twlev](https://ocean.weather.gov/estofs/estofs_surge_twlev.php)
792 [.php](https://ocean.weather.gov/estofs/estofs_surge_twlev.php)
- 793 NSF-NHERI. (2023). *Designsafe*. Retrieved 22.11.2023, from [https://](https://www.designsafe-ci.org/)
794 www.designsafe-ci.org/
- 795 Pachev, B., Arora, P., del Castillo-Negrete, C., Valseth, E., & Dawson, C. (2023).
796 A framework for flexible peak storm surge prediction. *Coastal Engineer-*
797 *ing, 186*, 104406. Retrieved from [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S0378383923001308)
798 [science/article/pii/S0378383923001308](https://www.sciencedirect.com/science/article/pii/S0378383923001308) doi: [https://doi.org/10.1016/](https://doi.org/10.1016/j.coastaleng.2023.104406)
799 [j.coastaleng.2023.104406](https://doi.org/10.1016/j.coastaleng.2023.104406)
- 800 Pe'eri, S. (2023). *Updated: Upgrade of the surge and tide operational forecast system*
801 *(stofs, formerly estofs) to version 1.1.0: Effective january 10, 2023* (Tech.
802 Rep.). NOS/Office of Coast Survey. Retrieved from [https://www.weather](https://www.weather.gov/media/notification/pdf2/scn22-108_stofs_v1.1.0_aab.pdf)
803 [.gov/media/notification/pdf2/scn22-108_stofs_v1.1.0_aab.pdf](https://www.weather.gov/media/notification/pdf2/scn22-108_stofs_v1.1.0_aab.pdf)
- 804 Pringle, W. J., Wirasaet, D., Roberts, K. J., & Westerink, J. J. (2021, February).
805 Global storm tide modeling with ADCIRC v55: unstructured mesh design and
806 performance. *Geoscientific Model Development, 14*(2), 1125–1145.
- 807 Pringle, W. J., Wirasaet, D., Suhardjo, A., Meixner, J., Westerink, J. J., Kennedy,
808 A. B., & Nong, S. (2018). Finite-element barotropic model for the indian and
809 western pacific oceans: Tidal model-data comparisons and sensitivities. *Ocean*
810 *Modelling, 129*, 13-38. Retrieved from [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S146350031830026X)
811 [science/article/pii/S146350031830026X](https://www.sciencedirect.com/science/article/pii/S146350031830026X) doi: [https://doi.org/10.1016/](https://doi.org/10.1016/j.ocemod.2018.07.003)
812 [j.ocemod.2018.07.003](https://doi.org/10.1016/j.ocemod.2018.07.003)
- 813 Ristic, B., Arulampalam, S., & Gordon, N. (2003). *Beyond the kalman filter: Parti-*
814 *cle filters for tracking applications*. Artech House.
- 815 Roelvink, J. A., & van Banning, G. (1995). Design and development of delft3d and
816 application to coastal morphodynamics. *Oceanographic Literature Review,*
817 *11*, 925. Retrieved from [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:127930966)
818 [127930966](https://api.semanticscholar.org/CorpusID:127930966)
- 819 Rougier, J., Brady, A., Bamber, J., Chuter, S., Royston, S., Vishwakarma, B. D.,
820 ... Ziegler, Y. (2023). The scope of the kalman filter for spatio-temporal ap-
821 plications in environmental science. *Environmetrics, 34*(1), e2773. Retrieved
822 from <https://onlinelibrary.wiley.com/doi/abs/10.1002/env.2773> doi:
823 <https://doi.org/10.1002/env.2773>
- 824 Schaffer, J., Timmermann, R., Arndt, J. E., Kristensen, S. S., Mayer, C.,
825 Morlighem, M., & Steinhage, D. (2016). A global, high-resolution data set
826 of ice sheet topography, cavity geometry, and ocean bathymetry. *Earth System*
827 *Science Data, 8*(2), 543–557. Retrieved from [https://essd.copernicus.org/](https://essd.copernicus.org/articles/8/543/2016/)
828 [articles/8/543/2016/](https://essd.copernicus.org/articles/8/543/2016/) doi: 10.5194/essd-8-543-2016
- 829 Seroka, G., Funakoshi, Y., Yang, Z., Moghimi, S., Myers, E., Pe'eri, S., ... Kaiser,
830 C. (2023, January). Upgrades to NOAA/NOS' surge and tide operational
831 forecast system (STOFS). In *103rd AMS annual meeting*. AMS.
- 832 Stammer, D., Ray, R. D., Andersen, O. B., Arbic, B. K., Bosch, W., Carrère, L.,
833 ... Yi, Y. (2014). Accuracy assessment of global barotropic ocean tide
834 models. *Reviews of Geophysics, 52*(3), 243-282. Retrieved from [https://](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014RG000450)
835 agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014RG000450 doi:

- 836 <https://doi.org/10.1002/2014RG000450>
- 837 Tiggeloven, T., Couasnon, A., van Straaten, C., Muis, S., & Ward, P. J. (2021, Au-
- 838 gust). Exploring deep learning capabilities for surge predictions in coastal ar-
- 839 eas. *Sci. Rep.*, *11*(1), 17224.
- 840 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ...
- 841 Polosukhin, I. (2017). Attention is all you need. *CoRR*, *abs/1706.03762*.
- 842 Retrieved from <http://arxiv.org/abs/1706.03762>
- 843 Verlaan, M., De Kleermaeker, S., & Buckman, L. (2015). Glossis: Global storm
- 844 surge forecasting and information system. In *Australasian coasts & ports*
- 845 *conference 2015: 22nd australasian coastal and ocean engineering confer-*
- 846 *ence and the 15th australasian port and harbour conference*. Auckland, New
- 847 Zealand. Retrieved from [https://search.informit.org/doi/10.3316/](https://search.informit.org/doi/10.3316/informit.703696922952912)
- 848 [informit.703696922952912](https://search.informit.org/doi/10.3316/informit.703696922952912)
- 849 Wang, P., & Bernier, N. B. (2023). Adding sea ice effects to a global opera-
- 850 tional model (nemo v3.6) for forecasting total water level: approach and
- 851 impact. *Geoscientific Model Development*, *16*(11), 3335–3354. Retrieved
- 852 from <https://gmd.copernicus.org/articles/16/3335/2023/> doi:
- 853 [10.5194/gmd-16-3335-2023](https://gmd.copernicus.org/articles/16/3335/2023/)
- 854 Wang, P., Bernier, N. B., Thompson, K. R., & Kodaira, T. (2021). Evaluation of
- 855 a global total water level model in the presence of radiational s2 tide. *Ocean*
- 856 *Modelling*, *168*, 101893. Retrieved from [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S1463500321001463)
- 857 [science/article/pii/S1463500321001463](https://www.sciencedirect.com/science/article/pii/S1463500321001463) doi: [https://doi.org/10.1016/](https://doi.org/10.1016/j.ocemod.2021.101893)
- 858 [j.ocemod.2021.101893](https://doi.org/10.1016/j.ocemod.2021.101893)
- 859 Westerink, J. J., Luettich, R. A., Feyen, J. C., Atkinson, J. H., Dawson, C., Roberts,
- 860 H. J., ... Pourtaheri, H. (2008, March). A basin- to Channel-Scale unstruc-
- 861 tured grid hurricane storm surge model applied to southern louisiana. *Mon.*
- 862 *Weather Rev.*, *136*(3), 833–864.
- 863 Willmott, C. J. (1981). On the validation of models. *Physical Geography*, *2*(2), 184-
- 864 194. Retrieved from <https://doi.org/10.1080/02723646.1981.10642213>
- 865 doi: [10.1080/02723646.1981.10642213](https://doi.org/10.1080/02723646.1981.10642213)
- 866 Wood, D. (2023). *Global storm and tide operational forecast system - global stofs*.
- 867 Notre Dame Computational Hydraulics Laboratory. Retrieved from [https://](https://dylnwood.github.io/GESTOFS-develop/)
- 868 dylnwood.github.io/GESTOFS-develop/
- 869 Xie, W., Xu, G., Zhang, H., & Dong, C. (2023). Developing a deep learning-
- 870 based storm surge forecasting model. *Ocean Modelling*, *182*, 102179. Re-
- 871 trieved from [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S1463500323000203)
- 872 [S1463500323000203](https://www.sciencedirect.com/science/article/pii/S1463500323000203) doi: <https://doi.org/10.1016/j.ocemod.2023.102179>
- 873 Zampieri, L., Arduini, G., Holland, M., Keeley, S. P. E., Mogensén, K., Shupe,
- 874 M. D., & Tietsche, S. (2023, May). A machine learning correction model of
- 875 the winter Clear-Sky temperature bias over the arctic sea ice in atmospheric
- 876 reanalyses. *Mon. Weather Rev.*, *151*(6), 1443–1458.