JGR | Machine Learning and Computation

RESEARCH ARTICLE
10.1029/2024JH000525

RTide: Automating the Tidal Response Method

Thomas Monahan[1], Tianning Tang[1], Stephen Roberts[1], and Thomas A. A. Adcock[1]

[1]Department of Engineering Science, University of Oxford, Oxford, UK

Key Points:

- Specially constrained neural network automates the response method, enabling analysis and prediction of nonstationary tidal processes
- A method is developed to explain any prediction and analyze the interactions between non-astronomical forcing and tidal processes
- RTide Python package enables full analysis and prediction of sea-levels under compound forcing using three lines of code

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

T. Monahan,
thomas.monahan@eng.ox.ac.uk

Citation:

Monahan, T., Tang, T., Roberts, S., & Adcock, T. A. A. (2025). RTide: Automating the tidal response method. *Journal of Geophysical Research: Machine Learning and Computation*, *2*, e2024JH000525. https://doi.org/10.1029/2024JH000525

Received 19 NOV 2024
Accepted 25 APR 2025

Author Contributions:

Conceptualization: Thomas Monahan
Data curation: Thomas Monahan
Formal analysis: Thomas Monahan
Funding acquisition: Thomas A. A. Adcock
Investigation: Thomas Monahan
Methodology: Thomas Monahan, Tianning Tang, Stephen Roberts, Thomas A. A. Adcock
Project administration: Thomas Monahan
Software: Thomas Monahan

Abstract Nonstationary tidal processes, such as tidal rivers and storm surge, present challenges for analysis and prediction because their inherent nonstationarity encumbers the use of standard tidal analysis tools like harmonic analysis. Moreover, specialized approaches impose problem-specific functional forms and rely on auxiliary data, limiting their applicability across different nonstationary tidal processes. Although Munk and Cartwright's tidal response method avoids these assumptions, its lack of automation has hindered broader application. Here, we develop a nonparametric, automated response-based analysis procedure. Our approach embeds a class of neural networks capable of representing any arbitrary Volterra series—the mathematical basis of the response method—within the classic framework. Our model facilitates the inclusion of meteorological and other non-tidal forcing. By explicitly accounting for nonstationarity, our method yields improved astronomical tidal estimates. We further devise a strategy to extract physical insights from the learned model, demonstrating its utility in studying the interaction and modulation of astronomical tides by external forcing. By taking a nonparametric approach, our framework enables the investigation of phenomena that heretofore could not be accounted for straightforwardly, as illustrated by several case studies on tide–surge interaction, riverine tides, and storm surge. These applications, and more, can be replicated with just three lines of code using the open-source Python package, RTide.

Plain Language Summary Accurate sea-level predictions are vital for coastal engineering, resilience, and understanding the combined effects of tides and other coastal processes on industries and ecosystems. Traditional methods for predicting tides often fall short when influenced by changing weather or other non-tidal factors like river outflow changes. Although the response method by Munk and Cartwright allows for the inclusion of such factors, it's difficult to automate and challenging to apply to study unknown physical processes without specialist expertise. Our new ML Response Framework combines this approach with a special class of neural network which is mathematically equivalent to the original method. This approach automates the response method and makes it easy to include non-tidal influences like weather, improves prediction accuracy, and can reveal how tides interact with other forces. The method is available in the easy-to-use RTide Python package, with case studies covering tide-surge effects, river tides, and nuisance flooding.

1. Introduction

The analysis and prediction of water levels driven by astronomical and other phenomena is important for countless scientific and engineering applications. As tides can become distorted in coastal and estuarine environments, we here define these processes by the forcing which gives rise to them. Tides are the weakly nonlinear response of sea-levels to the combined gravitational and radiational forcing from the moon and Sun. Radiational forcing refers to non-gravitational forcing such as surface pressure fluctuations from atmospheric tides. Tidally driven phenomena by extension are oceanic processes which exhibit coupled and weakly nonlinear responses to both tidal forcing and other external forces. Prediction of these processes is done using both numerical and empirical approaches which will be discussed below. In contrast, tools which can also be employed for the analysis of these phenomena are largely limited to tidal harmonic analysis (HA) (Parker, 2007) with notable exceptions for particular phenomena such as tidal rivers (Matte et al., 2014). While in most cases HA is sufficient for the study of tides, tidally driven phenomena such as tidal rivers and storm surges often violate the assumptions implicit within this formulation (Monahan et al., 2025). Application of HA in these contexts leads to biased estimates of tidal constituents and limited predictive power of the coupled process (Williams, Irazoqui Apecechea, et al., 2018). As will be shown, it is inappropriate to apply HA in these contexts, yet no general purpose method exists for both the analysis and prediction of these processes. Consequently, empirical approaches rely on concurrent observational data and thus struggle if such data is unavailable (e.g., predictions beyond short-term or if gauges are decommissioned). This manuscript addresses this deficiency by automating the tidal response

Supervision: Tianning Tang, Stephen Roberts, Thomas A. A. Adcock
Validation: Thomas Monahan
Visualization: Thomas Monahan
Writing – original draft: Thomas Monahan, Tianning Tang, Stephen Roberts, Thomas A. A. Adcock
Writing – review & editing: Thomas Monahan, Tianning Tang, Stephen Roberts, Thomas A. A. Adcock

method devised by Munk and Cartwright in the 1960s (Munk & Cartwright, 1966). The resultant method enables empirical analysis and prediction of any tidally driven process without requiring concurrent water-level measurements. To highlight the flexibility of our approach, all case studies in this manuscript were run using just three lines of code, without modification.

1.1. Literature Review

To contextualize our approach, we here look to briefly summarize the methods currently employed for analysis and prediction of tides and tidally driven phenomena.

The standard approach to predicting past and future tides, based on sea-level observations, is through harmonic analysis (HA) (Parker, 2007). As proposed by Darwin and Kelvin (Doodson, 1958), with refinements from Doodson (1921), HA describes the tides as the sum of oscillating tidal constituents whose amplitudes and phases are typically determined through least-squares fitting to the observed sea-level. The linear constituents described by HA exist at nearly identical frequencies to the astronomical forcing (Egbert & Ray, 2017). Shallow water effects in HA are handled by including nonlinear constituents whose frequencies are the sums or differences of linear constituents. This approach imposes strict assumptions about the data. In particular, conventional harmonic analysis assumes the signal to be stationary and that the oceanic response is coherent with the astronomical forcing which gives rise to it (Parker, 2007). This characteristic becomes problematic when HA is applied to observations contaminated with non-tidal forcing which can result in significant biasing of the derived constituents (Williams, Irazoqui Apecechea, et al., 2018). In addition to being a large source of error for operational surge forecasting, biased tidal estimates limit our ability to study both non-tidal and tidally driven processes as will be shown in the three subsequent case studies. Furthermore, evaluating tides within a harmonic framework limits the physical conclusions that can be drawn about the dynamics of the interactions of tides with non-gravitational forcing.

Machine learning methods have shown great promise in improving forecast accuracy over both conventional numerical and empirical approaches across many scientific applications (Brunton et al., 2020; Kratzert et al., 2019; Lam et al., 2023). Due to the great success of HA, applications of ML to pure tidal height forecasting are limited. The nonlinearity of tidal currents and their importance for tidal energy has justified the development of several ML approaches for tidal current forecasting confined to mainly short-term predictions (Monahan et al., 2023; Sarkar et al., 2016, 2018). Several works have developed ML approaches for storm surge forecasting (Chen et al., 2021; De Oliveira et al., 2009; Qin et al., 2023). Many approaches have been tried, however, the vast majority conform to a similar approach: based on water-level observations, have a model learn how to predict the next time-step recursively based on forecasted meteorological conditions. The variation in these methods comes from both different choices of input data and preprocessing as well as various model architectures (Qin et al., 2023). While these works demonstrate compelling reductions in error compared to standard numerical models, none are used operationally (Qin et al., 2023). This is a consequence of two main shortcomings. First, reliance on in-situ observations is problematic as gauges can go offline, or have variable accuracy due to biofouling and environmental conditions. The accuracy of the final model is thus fundamentally intertwined with the operational conditions of the gauge used to train the model. Furthermore, these approaches can not be used where observational data no longer exist. This is problematic as it prevents the application of these methods in ungauged regions (see Section 4.1). Second, as a consequence of their autoregressive and black-box construction, these methods have limited physical interpretability (Qin et al., 2023). This stems from working from observations rather than forcing. Here, the model is learning statistical correlations between these observables which describe their evolution. It is thus difficult to attribute the role of each input in producing a given forecast and model generalization is limited by past observed conditions.

Several specialized approaches have been developed for the analysis and prediction of nonstationary sea-levels. First, the familiar harmonic analysis formalism has been adapted for the prediction of riverine tides by replacing the conventionally static amplitudes and phases with functions of the river outflow and the diurnal tidal range from an adjacent location free from fluvial influence (Matte et al., 2014). This approach, termed "NS_Tide," derives its basis functions from a simplified fluvial model, and uses the regression to iteratively optimize coefficients for deviations in theory. This approach faces two key shortcomings; it requires (a) data from an adjacent reference station free from fluvial influence, and (b) it has only been developed for fluvial (using river outflow) and atmospheric pressure forcing (Matte et al., 2014). As such, even for tidal rivers, the approach loses

applicability if either an adjacent reference station or river outflow data is not available. Section 4.1 considers an application where both of these things are true. Furthermore, while "NS_Tide" can be effective for routine analysis and prediction of tidal rivers, it cannot be viewed as a general tool for analysis and prediction of nonstationary tidal processes. Other approaches, such as the species concordance method, similarly rely on uncontaminated reference series and thus face similar limitations (George & Simon, 1984). Many approaches have been developed for analysis of tidal processes, but offer no predictive value and are thus out of the scope of this discussion (Jay & Kukulka, 2003; Lobo et al., 2024; Pan et al., 2018).

The complex modifications of HA to accommodate nonstationary effects are an artifice arising from restricting the solutions to sums of discrete harmonic functions of time. In contrast, the Response method provides a framework for computing the weakly nonlinear oceanic response to a set of input forcing (Munk & Cartwright, 1966). This approach can be viewed as learning a time-invariant operator from the forcing to the associated sea-level (described at length in Section 2.1). Originally, this was just done just from the tidal input potential but was later applied to meteorological forcing (Cartwright, 1968). Under this construction, the oceanic response is explicitly separated out from the input forcing, providing the ability to attribute the contribution from each input force. The learned model can therefore be used to identify the interaction of tides with these forces and can be used to predict the combined response of any tidally driven phenomenon. Many works have found response approaches to yield modest improvements over HA for pure tidal prediction (Cartwright & Ray, 1990; Munk & Cartwright, 1966). However, applications of the approach to predict nonstationary tidally driven phenomena have been limited outside of Cartwright, 1968. This is a consequence of the approach being challenging to automate as nonstationary tidal processes exhibit irregular and nonlinear response functions (see Section 2.1.1) (Cartwright, 1968; Groves & Reynolds, 1975; Munk & Cartwright, 1966). The remainder of this manuscript addresses this limitation and develops a general-purpose analysis and prediction tool for tidal processes.

2. Methods

2.1. Response Method

The Response method seeks to describe the time-invariant oceanic response at a point to a series of impulses $V(t)$. The time-series representing the tidal input potential $V(t)$, can be computed at any time and location based on the known orbits of the Moon, and Sun. A key contribution of the original response method was the recognition that the tide at a given point was not merely driven by the instantaneous equilibrium potential, but rather a combination of its past, present, and future values (Munk & Cartwright, 1966). Hence, the Response method expands the gravitational tidal potential using spherical harmonics such that

$$V(\theta,\lambda;t) = g\sum_{n=0}^{n}\sum_{m=0}^{n}\left[a_n^m(t)U_n^m(\theta,\lambda) + b_n^m(t)V_n^m(\theta,\lambda)\right], \tag{1}$$

where g is the gravitational constant, θ is the Geographical Colatitude, λ the East longitude, U_n^m and iV_n^m are complex spherical harmonics of degree n and order m, and a_n^m and b_n^m are the associated complex values of the global tide function $c_n^m(t)$, yielding $V(\theta,\lambda;t) \in \mathbb{C}$. The spherical harmonics converge rapidly and are defined so that the tidal species of interest is governed by the selection of m. Normalization of the spherical harmonics is consistent with quantum mechanics (Donini et al., 1977). A table containing the frequency bands for each species is provided in (Munk & Cartwright, 1966).

The linear response $\hat{\zeta}(t)$ is given by the convolution of response weights $w_n^m = u_n^m + iv_n^m$ and the global tide function $c_n^m = a_n^m + ib_n^m$ at various time-lags τ such that

$$\hat{\zeta}(t) = \sum_{m,n}\sum_{s}\left[u_n^m(s)a_n^m(t-\tau_s) + v_n^m(s)b_n^m(t-\tau_s)\right], \tag{2}$$

with the response weights determined by least squares and $\hat{\zeta} \in \mathbb{R}$. Higher order terms are computed using either a sequential or a lumped analysis procedure by forming products of the linear terms. The x^{th} order response, with degree and order m,n, takes the form of

$$x^{th}\text{order response} = \sum_i \dots \sum_x \sum_s \dots \sum_{s'} w(i,\dots,x,s,\dots,s')\,(c(t-\tau_s))(\dots)(\dots)(c(t-\tau_{s'})), \tag{3}$$

with the number of terms increasing exponentially with the order of the interactions. Typically, higher order terms are limited to only one lag for computational efficiency. Due to the lack of orthogonality of the basis functions for the standard response method, much care is required when performing a sequential analysis. The time lag, τ_s, is typically 2 days and s is the maximum lag with the understanding that convolutions are carried out using $\tau = 0, \Delta\tau, 2\Delta\tau, \dots, s\Delta\tau$. An extended discussion of the choice of time lags and their implications on the associated response is given in Section 2.2.3.

The response formalism can also be used to compute the oceanic response to non-gravitational forcing. This is accomplished by adding additional input functions into the analysis. Here, the term input function refers to the type of forcing included in a given analysis. Individual inputs are given by time-lags of the associated input functions. Munk and Cartwright propose an additional input function termed the "radiational flux." This is used to account for both daily wind and pressure fluctuations and seasonal changes in oceanic temperature. Unlike gravitational forcing, they assume the Earth is "opaque" to radiational forcing. As such, at night the radiational forcing $\mathcal{R} = 0$, and during the day is proportional to the local zenith angle α of the sun with $\mathcal{R} = S \cdot \cos\alpha$. Here S can be thought of as the incident radiant energy. This simplified model is then expanded into spherical harmonics as with the gravitational inputs. Our testing showed inclusion of radiational input functions can be useful for analyzing shortened reference series (less than 1 year) and accounting for seasonal effects. However, these inputs are not necessarily physical but rather approximations of a number of different non-gravitational effects. When analyzing the learned model, interactions with radiational forcing cannot be considered physical.

A complete description of the Gravitational and Radiational input functions can be found in their original paper (Munk & Cartwright, 1966). Additionally, code to compute both gravitational and radiational input functions is provided in the RTide python package.

2.1.1. Challenges of Conventional Response Analysis

While the response method possesses several advantages over HA, it carries with it what is referred to as the burden of realistic input function selection (Munk & Cartwright, 1966). This language is in someways a misnomer as the input forcing functions for a standard response analysis are readily computed from the derivations given in (Munk & Cartwright, 1966). We argue the difficulty in a response analysis is not introduced by the input functions themselves, but instead, the need to specify the analytical form of their interactions *a priori*. Here, we refer to these relationships as *interaction kernels*, as this description is consistent with the broader literature on Volterra series. Similar to the selection of harmonic constituents in the harmonic method, the standard response approach requires the definition of these interaction kernels beforehand. This requirement is in some ways more restrictive than the selection of constituents as these interactions becomes increasingly complex and consequently difficult to prescribe in strongly nonlinear regions. Further difficulty is introduced if one wishes to include other forcing inputs as is shown in Cartwright, 1968. Indeed, as will be shown in Section 4.1, the nature of these interactions for many multivariate inputs is not known. This is made more difficult to debug as the response weights generally lose meaning for such inputs (Cartwright, 1968). As such, the standard response approach often requires expert trial and error which has severely limited its automated application.

The method proposed here-in alleviates this burden by simultaneously inferring the weights and the structure of the interaction kernels using a data-driven approach. This method can learn high-order (greater than second order) nonlinear interactions associated with multivariate inputs and fully automates the response analysis. Subsequent sections illustrate how this non-parametric formulation is both useful for practical sea level prediction of complex tidal phenomena and studying how the sea level responds to such forcing.

2.2. ML Response Framework

In this section, we introduce our proposed ML Response Framework. The basis of our approach is the integration of a specific class of three-layer neural networks into the tidal impulse-response theory. That is, rather than specify a parametric form for the interaction kernels before an analysis and infer the response weights using least-squares, as in the original method, we instead use the networks to infer both the interactions between inputs and

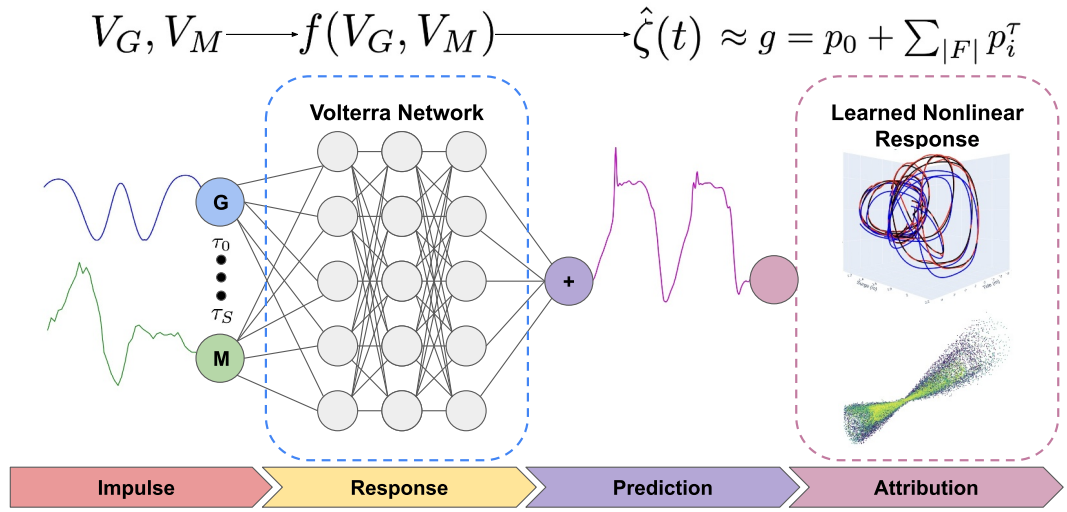

Figure 1. Schematic of the proposed ML Response framework. The input consists of the lagged Gravitational (V_G) and Multivariate (V_M) input functions. Here multivariate can be anything from the "radiative flux" proposed in (Munk & Cartwright, 1966) to river levels or atmospheric variables. The oceanic response function $f(V_G, V_M)$, equivalent to Equations 2 and 3, is learned by a three-layer neural network (Volterra network) and used to predict either total sea-level $\hat{\zeta}$ or the astronomical tide $\hat{\zeta_T}$ at any time instance t. Learned models can then be taken apart and studied by transforming into an equivalent linear model g using SHapley Additive feature exPlainers which quantify the contribution of each feature p_i^{τ} to the final prediction (Lundberg & Lee, 2017) outlined in Section 2.3.

the associated weights simultaneously and directly from the data. The motivation for this procedure lies in the provable equivalence between the Volterra series—the mathematical basis of Munk and Cartwright's response method—and the aforementioned constrained three-layer-perceptrons (Wray & Green, 1994). This relationship makes the approximation of any arbitrary combination of Volterra kernels possible (Wray & Green, 1994). Thus, our framework can infer the appropriate Volterra kernel structure directly from the available data and efficiently represent higher order nonlinearities (e.g., 3rd order and above) if required (Andersson et al., 2019).

The response prediction, defined by Equation 2, can be replaced by

$$\hat{\zeta}(t) = f(V_0(t-\tau_0), V_0(t-\tau_1), \ldots, V_0(t-\tau_S), \ldots, V_X(t-\tau_0), V_X(t-\tau_1), \ldots, V_X(t-\tau_S)), \tag{4}$$

where $f(\cdot)$ is the NN approximation of the time-invariant impulse-response function and V_i are the set of all input functions at the associated time lags. It should be noted that for input functions expanded in spherical harmonics, inputs of differing degrees and orders are provided as separate inputs, and thus considered separate input functions, as in the original formalism. In contrast with the original method, the data does not need to be bandpassed beforehand as all species and interactions are inferred simultaneously. The underlying mathematical connection to the Volterra series allows the networks, as deployed herein, to simultaneously learn both the linear time invariant impulse-response, as well as the associated nonlinear interactions, which are challenging for the original rigid convolution method. For the analysis presented in this paper, the network weights w are computed using gradient descent. This non-parametric procedure retains the physical advantages and realism of the original response method but overcomes the limitations of the parametric original method by allowing the model to simultaneously learn the structure of the input function interactions and weights directly from the data. The simultaneous inference of feature weights and interactions serves as a natural regularizer, thereby avoiding some of the variance-bias trade-off imposed by least squares. Critically, when paired with the training and hyperparameter optimization procedures presented in Sections 2.2.5 and 2.2.4, this framing reduces the need for expert tuning and relaxes the need for the explicit definition and expansion of input functions. The following section outlines the implementation of the method used in this work and in the RTide Python package, and describes its relation to the original response method. Figure 1 provides a schematic of our approach.

2.2.1. Network Architecture

We use a class of three layer feed forward neural networks (NN) termed "Volterra Neural Networks" due to their ability to represent any arbitrary Volterra series (Marmarelis & Zhao, 1997). The basic NN architecture has three layers with size equal to the number of inputs. If we have X input functions and S time-lags then the size of each layer is simply $X \times S$. We find little sensitivity to the selection of activation functions, but utilize the sigmoidal activation function, tanh, which allows for the indirect estimation of Volterra kernels from the learned network (Marmarelis & Zhao, 1997). The final layer consists of a simple adder which sums the outputs directly from the previous layer. The choice of this architecture reflects our objective to produce a non-parametric approach that is mathematically equivalent to the original method, but critically removes the need to select interaction terms *a priori*. As a consequence of this equivalence, these networks are generally robust and can be applied universally. In our testing (results not included in this manuscript), we have found that structuring networks in accordance with known physical interactions between input functions (e.g., tide-surge interactions) can improve the resolution of the underlying dynamics. We believe the development of such inductive biases to be a promising avenue for future work.

A key modification in our framework is the use of a "deep" rather than "wide" network. A wide network is comprised of a single hidden layer and is equivalent to a generalized linear model. This construction yields excellent memorization capabilities, but will often struggle with generalization (Cheng et al., 2016). As such, it is easy to overfit these models which will lead to poor performance on unseen data. Conversely, deep networks, with multiple hidden layers, demonstrate superior generalization as they can form low-dimensional embeddings from the input space. A wide model is the appropriate choice when applied to low-noise data, and in predominantly linear tidal regimes. However, we find that the improved generalization of deep networks is beneficial for studying complex tidal regions with non-tidal forcing.

Typical training times are approximately 3 min when applied to 1-year of hourly data (8760 training points) using $X = 10$ input functions with $S = 7$ time-lags ($3 \times (10 \times 7)^2 + 3 \times 10 \times 7 + 10 \times 7 + 1 = 14981$ trainable parameters) using a Tesla K80 GPU. While this reflects the total number of trainable parameters, the regularization of the model reduces the effective number of degrees of freedom drastically. As such, the data helps to inform the complexity of the model which we believe to be a major advantage. We note that training time scales with both the number of measurements and the number of input functions and time lags included.

2.2.2. Input Forcing

This work defines standard analysis as any analysis which makes use of just the gravitational input functions defined in (Munk & Cartwright, 1966). These are implemented exactly as in the original work up to degree 3 for the gravitational inputs. Due to the self-regularization of our non-parametric approach, it was found useful to include all available degrees and orders of spherical harmonics for standard analysis. Thus, when combined with hyper parameter optimization, standard analysis can be fully automated.

As noted in (Munk & Cartwright, 1966), inclusion of the "radiational flux" can help combat the variable nature of seasonal and annual weather fluctuations. Depending on the location, and the consistency of this seasonal variability, inclusion of such inputs can either help or hurt long-term prediction skill. For Sections 3.1 and 4.1 we do not utilize radiational input functions as we wish to isolate and obtain physically meaningful insights regarding the oceanic response to different forcing. Section 4.2 briefly evaluates the impact of including Radiational inputs. Radiational inputs can be included in the RTide package by setting "radiational = True" in the "Prepare_Inputs" function. Due to their quasi-physical nature, the learned response should also be regarded as quasi-physical.

The flexibility of our framework enables the inclusion of any multivariate input forcing which can be represented as a time-series in addition to the original gravitational and radiational input functions given in (Munk & Cartwright, 1966). In this work, we also include what we term multivariate input functions. We define multivariate inputs as any forcing that is not gravitational. A consequence of our non-parametric framework is that we can relax the requirement for input function definition. As an example, in (Cartwright, 1968), the pressure gradient is obtained by fitting coefficients of a spatial Taylor series to meteorological observations at 8 reference stations. Our testing showed that our method could obtain similar performance when trained using the meteorological variables as direct input. In this way, our non-parametric formulation does not impose approximate physical assumptions such as the linearization of wind stress used in (Cartwright, 1968). Rather, our approach infers these

relationships directly from the data. This characteristic can thus be used to inform the design of such low order models and is described in Section 4.1. This capability makes the testing of new input functions extremely easy and the integration of other data streams trivial using only one line of code in RTide.

2.2.3. Lag Considerations and the Credo of Smoothness

The selection of which time lags and spacing to use, often referred to as *embedding parameters*, is critical to the success of a Response analysis. In their original paper, Munk and Cartwright propose a uniform (evenly spaced) time lag of $\Delta\tau = 2$ days based on empirical analysis of the Honolulu tidal record (Munk & Cartwright, 1966). Here, we refer to the selection of evenly spaced lags as uniform embedding. They further truncate their response analysis after $S = 3$ lags following the observations that the tidal admittance $Z(f)$ shows no indication of "wiggles" with period less than $1/6$ cycles per day. This selection of a maximum lag results in a smoothing of any portion of the signal with period less than $1/6$ cycles per day which they termed the "credo of smoothness." While our testing confirms these observations, we also found that the inclusion of additional non-uniformly spaced lags improved performance (See Table S1 in Supporting Information S1). We remark that while this approach has improved performance for our non-parametric method, the non-uniform lags may introduce complications when performing spectral analysis on the classic method.

We believe the justification for this modification comes from the dynamics of the tidal system. It is well understood that uniform embedding can yield poor results when applied to systems with multiple strong periodicities across disparate timescales (Judd & Mees, 1998). Non-uniform embedding provides the opportunity to characterize systems with fast-slow dynamics (Tan et al., 2023). While numerous data-driven approaches have been developed for the automated selection of non-uniform time embeddings, tidal flows provide a unique circumstance in which the characteristic periods are well understood. We utilize a non-uniform embedding scheme based on a subset of the natural tidal frequencies. The associated time lags τ_i are given by the quarter period heuristic which yields a maximally convex embedding in the reconstruction space if the characteristic frequencies can be appropriately identified (Fowler & Kember, 1993). We select S positive and negative time lags τ_i such that

$$\tau_i = \frac{2\pi}{4\omega_i} \tag{5}$$

where ω_i are the constituent frequencies that satisfy a Rayleigh criterion of 1 based on the length of the training data. While marginal improvements are garnered by the inclusion of additional lags, we find it sufficient to truncate our analysis to only include the constituents "resolvable" from a 7-day record with sampling rate 1 hr (M2, K1, M4, M6, M3, M8, 3MK7, 2SK5, 2MK5). We note that this embedding yields a maximum lag of ≈ 6.25 days and is therefore consistent with the original credo of smoothness. In addition to capturing dynamics on multiple timescales, this embedding procedure introduces a degree of explainability consistent with the classical harmonic method.

As proposed in (Woodworth & Vassie, 2022), we have experimented with a symmetrical analysis scheme that includes an equal number of positive and negative lags. We found this procedure to improve model performance (See Table S1 in Supporting Information S1). While leading time lags appear to violate causality, we propose the following explanation based on the method of characteristics. As a consequence of selecting a finite number of lags in our embedding procedure, we constrain our lag embedding to consider the first τ_i negative lags. Suppose we have an arbitrary signal whose characteristic period is not included in this procedure. If we consider the tidal signal to be a series of waves observed at a stationary location at time t governed by a single PDE, then the signals "originating" from each time lag can be considered to act along the Spatiotemporal-Characteristic of the tidal system. The inclusion of positive lags, therefore has a simple explanation; positive lags correspond to signals acting along the same characteristic, but intersecting a τ_i in the future (whether in time or space).

Applications that include non-astronomical forcing in this manuscript only consider negative and real-time lags. We note that the optimization of these lags warrants future work, and can be problem dependent, as the time scales of the associated interactions can differ significantly from those of the tidal forcing. The learned models can be used for this optimization by inspecting the feature importance of the individual time-lags of the multivariate forcing and tweaking the parameters accordingly.

2.2.4. Hyperparameters and Tuning

While the selection of input functions and interaction kernels has now been automated, the choice of hyperparameters plays an important role in the success of an RTide analysis. Here, hyperparameters refer to the parameters which govern the training characteristics of the model (training duration or "epochs," learning rate, and regularization strength). Reported values as well as general recommendations for hyperparameters are included in the RTide package. In our testing, we found the default hyperparameters included in the RTide package to yield state-of-the-art performance and generalization (balanced performance out of sample) across a broad range of problems of varying complexity and non-gravitational forcing. Indeed all case studies presented here have been produced using these standard parameters. However, to allow users to achieve maximal performance an adaptive hyperparameter selection scheme can be adopted. We elect to use the Keras Hyperband tuner due to its rapid convergence and strong performance NN architectures (Rala Cordeiro et al., 2021). The optimization procedure is carried out in a tournament style and is accelerated through optimal resource allocation and early stopping (Li et al., 2017). This procedure is carried out during the training procedure and repeated until a user defined threshold (either iterations, time, convergence, etc.) is met.

2.2.5. Training Procedure

For problems where data is abundant (e.g., ≥30 days of data with sampling rate ≤1 day), a standard training procedure is sufficient. That is, after weight initialization, the entire network is trained end-to-end using the computed input functions at the associated time lags as direct input. The input training time-series is randomly shuffled and automatically partitioned into a train and validation set. By default, this is set to 95% and 5% of the data respectively. This validation set is used throughout the network training to monitor the model's generalization capabilities. Inspection of the training set and validation set loss curves is helpful for diagnosing overfitting and is shown by default when performing RTide analysis. An example of this is shown in Figure S2 in Supporting Information S1.

Unlike conventional response analysis in which the interaction terms must be defined beforehand, the interactions between inputs and associated weights are simultaneously inferred through the training procedure using gradient descent. This allows the data to determine the structure of the model thereby acting as implicit regularization. Using the RTide implementation this allows for response analysis to be carried out using only three lines of code. Gradient descent simply looks to minimize the provided cost function. Hence, additional modifications could be applied to the loss function depending on the user's objectives. For example, minimizing the absolute error in water-level prediction will not always yield the best tidal estimate and vice versa. The results reported in this work only use the mean absolute error (MAE) as the chosen loss function. We note that early testing has shown the Surface Similarity Parameter (Wedler et al., 2022) can yield more accurate tidal estimates in some scenarios and is implemented in RTide. Unlike Mean Squared Error, MAE places less emphasis on outliers and is thus more robust to nonstationary artifacts. Dynamic tuning of the learning rate by the ADAM optimizer further reduces sensitivity (Kingma & Ba, 2014). We have also experimented with both L1 and L2 regularization of the network weights. Both of these have proved to be beneficial for reducing overfitting and reducing the effective number of degrees of freedom of the model. Regularization strength parameters are included in the hyperparameter tuning described in the previous section.

2.3. Model Interpretability – SHAP Feature Importance

Despite the equivalence of the Volterra network to the original analytical response formalism, the learned network itself is not interpretable. As such, we propose a method for transforming the network into an equivalent linear model. This approach allows us to attribute the contribution of each input function to any prediction thus allowing for conclusions to be made regarding the learned dynamics.

Shapley Additive exPlanations (SHAP) is a model agnostic framework for developing machine learning "explainer models" based on Shapley values from game theory (Lundberg & Lee, 2017). This framework allows us to transform the learned response model f into a simplified additive explanation model g defined by the linear combinations of binary variables p_k such that $f \approx g = p_0 + \sum_{|F|} p_k$ where F is the set of all input features and $|F|$ is the number of input features. For a given prediction $f(t)$, the associated Shapley values $p_i^\tau(t)$ for each input function i and time lag τ provide an allocation of credit among input features (e.g., the sum of SHAP values is

equal to the expected value of the model) (Lundberg et al., 2018). The Shapley value $p_i^\tau(t)$ of each feature can therefore be interpreted as the impact of each input function at the associated time lag on the oceanic response.

For each input feature, a model is constructed with the feature included and a separate model is built with the feature withheld. The models are then compared across the possible values of the input feature. To combat potential multicolinearity, SHAP computes these differences across all possible subsets of feature inputs $R \subseteq F/\left\{{}_i^\tau\right\}$. The final SHAP value p_i^τ for a given input feature is therefore determined by the weighted average of all computed differences such that:

$$p_i^\tau(t) = \sum_{R \in F/\left\{{}_i^\tau\right\}} \frac{|R|!(|F| - |R| - 1)!}{|F|!}\left[f_{R\cup\left\{{}_i^\tau\right\}}(t_{R\cup\left\{{}_i^\tau\right\}}) - f(t_R)\right] \tag{6}$$

Shapley values satisfy two important properties:

1. Local Accuracy: For all predictions, the sum of attributed values $p_i^\tau(t)$ will equal the difference between the expected value $E[f(t)]$ for the entire training set and the observed value $f(t)$.
2. Consistency: Regardless of other features present, if a feature is more important to a model than another, then the value attributed to that feature will be higher.

The combination of these properties makes it possible to explain why the learned response model makes its prediction and thus attributes physical causes of variability. Critically this analysis also allows for the comparison of feature importance between locations, times, and in the presence of features with partial coherence. Indeed, implicit within this is quantifying the interactions of different inputs as shown in Section 4.1. This is important as it can shed light on how non-gravitational forcing modulates the astronomical tide. A complete treatment of the SHAP algorithm including proofs of the two aforementioned properties can be found in (Lundberg & Lee, 2017). SHAP values in this work are computed using the SHAP python package (github.com/shap/shap).

2.4. Harmonic Analysis

The Harmonic Analysis described in this paper uses a Python adaptation of the MATLAB Harmonic Analysis program UTide (Codiga, 2011) with a Rayleigh threshold $R = 1$. Analysis is conducted using ordinary least squares (OLS) unless explicitly stated. A complete description of the harmonic analysis theory can also be found in (Parker, 2007).

3. Simulated Results

3.1. Nonlinear Impulse Response Discovery

While RTide can be employed for conventional tidal analysis problems, it is important to first draw a distinction between the parametric regression based formulation of conventional tidal analysis and our non-parametric formulation. We stress that the parametric formulation has the distinct advantage of being both fully interpretable and compact, and is thus the preferred method for well understood tidal regimes (e.g., longstanding tide-gauges, tides with minimal meteorological forcing, river tides when river outflow measurements are available, etc.). However, the same characteristics which make parametric approaches useful for particular cases are also severely limiting when we look to apply them to study new and more complex phenomena. Here we encounter a paradoxical situation: how do we look to study new tidal phenomena with methods which require understanding how these phenomena interact in the first place? We believe RTide closes this gap, by allowing researchers to explore these complex interactions without requiring a fixed regression formulation. To test the proposed approach, we generate 2-year of synthetic sea-level observations given by the superposition of a pure tidal signal with a surge and nonlinear tide-surge interaction signal, such that the total observed sea-level is given by:

$$\zeta = \zeta_T + \zeta_S + B * \zeta_S\left(\zeta_T'' + \zeta_T'\right), \tag{7}$$

where ζ_T and ζ_S are the independent tide and surge signals, with tide-surge interaction defined by a constant B, and the first and second derivatives of the tidal signal ζ_T' and ζ_T''. The simulated tide ζ_T, sampled every 15 min (70080 observations), is composed of 59 harmonics originally derived from 1-year of tide-gauge measurements at

Eastport, Maine. The signal is largely semi-diurnal with an amplitude of approximately 3 m. We approximate a surge using the formula outlined in Cartwright's original work on surges (Cartwright, 1968) such that $\zeta_S = A*(1 + \cos(2\pi t/l))$ where $6 \leq l \leq 96$ (hours) is the duration of the surge, defined from $0 \leq t \leq l$ and $0.1 \leq A \leq 2.1$ (m) is the amplitude of the surge. The nonlinear tide-surge interaction term is defined by the product of the pure surge with the first and second derivatives of the tidal signal respectively. This simplified model, can approximate the well documented behavior of tide-surge interactions in which surge effects are amplified when they coincide with a low and rising tide and damped at high and falling tides (Doodson, 1929; Keers, 1968; Prandle & Wolf, 1978). It should be noted that this is a low-order approximation developed from observations in order to validate synthetic models of tide-surge interaction and is not a true representation of the physics. $B = 1$ is chosen such that the tide-surge interaction never exceeds more than 50% of the actual surge (Idier et al., 2012). We elect to use this simplified tide-surge formulation as it is easily interpretable within the time domain, and illustrates the importance of our approach even if the surge signal was perfectly predicted. Simulated surges are randomly generated using the defined parameter bounds and randomly added to the synthetic tidal signal so that a surge event occurs within each 2-week window. In total, there are 52 surges of variable magnitude and duration applied to the tidal signal. White Gaussian Noise is added with a SNR of 60 to represent more realistic conditions based on analysis of the original Eastport tide-gauge data. Visualizations of the simulated surge data can be found in Supplementary Information Figure S1 in Supporting Information S1. Section 4.2 shows how RTide can be used to learn the nonlinear response to meteorological forcing for the more challenging problem when the surge is generated by multiple meteorological processes.

Here, we suppose that the surge is produced from fluctuations in atmospheric pressure through the inverse barometer effect such that $\zeta_S = -0.01 \times P$ where ζ_S is in meters and P is in mbar (Fu & Pihos, 1994). In this case-study, let us also assume the atmospheric pressure is accurately forecasted and can thus be used as an input function into RTide. Hence, both the surge and the nonlinear relationship between tide and surge is unknown a priori. Figure 2 compares the performance of RTide and HA when used to recover the tidal signal, storm surge (IB effect), and nonlinear tide-surge interaction. Both models are trained on the first 1.5 years (550 days) of synthetic data.

Before considering any non-tidal effects, we first evaluate the performance of conventional least-squares (OLS) HA with RTide. Figure 2 Panel (A) shows that the process of simultaneous tide, surge, and tide-surge estimation results in the more accurate recovery of the true tidal signal by RTide. The RTide predicted tidal signal is simply obtained by passing a vector of zeros for the normalized barometric pressure input functions. This improvement is reflected by the reduction in mean absolute error (MAE) in the derived tidal signal from 0.054 m for HA to 0.009 m for RTide. We note that this comparison is conducted using HA with OLS, however, we have also compared the performance to modified HA using iteratively reweighed least squares (IRLS). While IRLS improves performance over OLS for contaminated signals, the RTide tidal prediction still yields a reduction in MAE from 0.010 to 0.009 m.

Due to the fact that our system is separable, one may assume that the linear inverse barometer effect could be estimated directly by simply regressing the atmospheric pressure P against the non-tidal residual (surge + tide-surge interaction). This procedure assumes the remaining variability after removing the harmonic or response computed tidal signal is Gaussian distributed noise. Figure 2 Panel(s) B compare the HA and RTide estimated inverse barometer effect (IBE) by simple linear regression of the non-tidal residual with the atmospheric pressure. The impact of the biasing in Panel A is clearly shown. While RTide recovers the exact IBE of -1.002 cm/mbar, the HA obtains an IBE of -0.975 cm/mbar. While 2.5% error may seem minor, these errors compound with the tidal errors and any residual noise in the observations when estimating the tide-surge interaction.

Figure 2 Panel(s) (C) shows how the biasing of the derived HA tidal signal significantly degrades the accuracy of the residual tide-surge interaction. Here, we are comparing the learned and predicted tide-surge interaction from RTide with the tide-surge interaction obtained by subtraction of the HA tidal prediction. It can be seen that the RTide learned and predicted tide-surge interaction converges to the true system, as is evidenced by R^2s of 0.999 and 0.923 on the train and test sets respectively. The degradation in the RTide estimated tide-surge interaction between train and test-set is largely from errors at low tide-surge values in which tidal errors also contribute. In comparison, the HA residual tide-surge interaction shows an R^2 of 0.883 for the train set indicating that the biasing of the tidal prediction results in a degradation of the true nonlinear system once subtracted. Since this is only a residual and not a prediction, an equivalent test set prediction is not possible. The degradation of the HA R^2

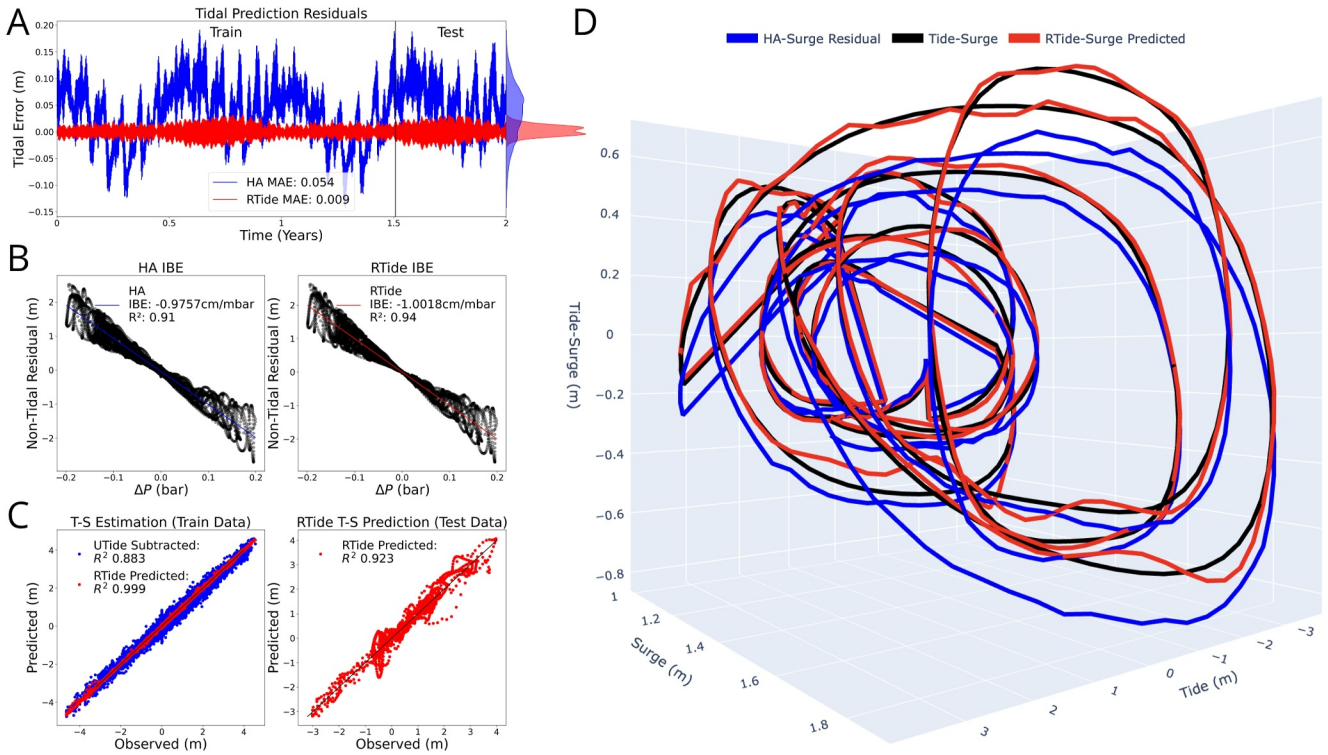

Figure 2. Comparisons of RTide and HA when used to recover both the nonlinear tide-surge interaction and the tidal signal. Panel (a) shows the error in tidal predictions for HA and RTide when compared to the true tidal signal. Panel (b) shows the HA and RTide estimated inverse barometer effect (IBE) using a linear regression of the change in pressure ΔP and non-tidal residual. Panel (c) shows comparisons of the HA tide-surge residual and the RTide tide-surge prediction. Note, since the HA is only a residual, it cannot be used for prediction on the test data. Panel (d) shows HA residuals and response predictions of the tide-surge response for different values of tide and surge. Note, Panel (d) is truncated between surge values of 1 and 2.2 m for legibility. RTide uses surge values as an additional input function.

is primarily due to errors at low surge values but is also a function of the noise of the data. Given the low amount of Gaussian noise applied, this represents a best-case scenario. As shown, even infrequent non-stationary contamination can yield systematic tidal estimation error. This is further confirmed by Figure 2 Panel (D) in which the RTide demonstrates excellent agreement with the actual nonlinear system and the HA is severely biased. These results indicate that RTide can learn the true nonlinear impulse-response relationships directly from the data. This result is important for application on real-data where studying these relationships using conventional parametric methods is challenging. Furthermore, we have demonstrated the shortcoming of using standard harmonic methods to study systems contaminated by non tidal forcing. We stress that any analysis conducted on the biased residual will result in similarly biased conclusions.

Similar results are observed when comparing the accuracy of the derived harmonic constituents. The error of the k^{th} constituent in terms of the quadrature amplitudes is estimated using the RMS (Stammer et al., 2014), defined as

$$\mathrm{RMS}_k = \sqrt{\overline{[(A_{\mathrm{model}}\sin(\omega_k t) + B_{\mathrm{model}}\cos(\omega_k t)) - (A_{\mathrm{true}}\sin(\omega_k t) + B_{\mathrm{true}}\cos(\omega_k t))]^2}}. \quad (8)$$

This metric provides a balanced measure of the derived harmonic constituent error by penalizing both amplitude and phase proportionally. Results are provided in Table S2 in Supporting Information S1. It can be seen that by including the atmospheric pressure component as an additional input into RTide, a lower RMS error of the tidal constituents is obtained for a majority of constituents. An estimate of the total error across all constituents is given by the root-sum-of-squares (RSS), where $\mathrm{RSS} = \sqrt{\sum_{k=1}^{n}(\mathrm{RMS}_k)^2}$. The RTide derived tidal signal shows a lower RSS over both OLS and IRLS harmonic analysis methods. While the improvements over IRLS are only slight for this example, it is expected that these disparities will grow in the presence of a continuous noise spectrum.

Regardless, we emphasize that the optimal choice of method depends on the objectives of the researcher and that an optimal analysis may result from using both RTide and HA.

4. Real Data

Having shown the capabilities of RTide when applied to simulated data, we go on to demonstrate its ability on more complex problems using real-world data. We note that our aim in the present paper is to demonstrate the applicability of the approach to these problems. Hence, sites selected for these case studies were done so as to be maximally representative, though the method has been tried—and validated—on multiple locations. Future work will examine these in greater depth and refine the analysis from that presented here.

4.1. Tidal Rivers

Sitting at the boundary between hydrology and physical oceanography, tidal rivers represent a unique challenge in tidal modeling. The confluence of river outflow, variable bathymetry, and sediment input reduce the performance of conventional harmonic and numerical methods (Hoitink & Jay, 2016). Despite this, modeling these regions has significant implications for the global water budget, and in regional resource management and flood mitigation efforts (Wells, 1995). We demonstrate the improved ability of RTide to capture this behavior and to integrate multivariate data sources using data from beneath the Chester Weir, near Liverpool, UK for the year 2020. As a tidal river, the observed water levels at this location are a consequence of the nonlinear interaction between the incoming tide and the outflow of the River Dee.

While non-stationary HA methods have been developed for the prediction of tidal rivers, they require the presence of a secondary reference station free from non-stationary contamination which significantly reduces their global applicability (Gan et al., 2019; George & Simon, 1984; Matte et al., 2014). Additionally, the basis functions employed in NS_TIDE (Matte et al., 2014) require river outflow measurements, and not river water-levels. For many locations in the UK and globally, river outflow measurements are unavailable. To combat this, we use upstream water levels as direct input into our model, and use RTide to approximate the localized sea-level response. While river levels are not as accurate as river outflow, they are much cheaper to obtain and can even be extracted via satellite data. As such, the proposed approach removes the current requirement of outflow data and can be applied globally. It is in theory possible to define a basis function for NS_TIDE that could take upstream river levels as input. However, to the best of our knowledge, this has not been done in the literature and would require significant expertise input. Future work will look to analyze the learned RTide model to produce an approximate basis function that can be integrated into the NS_Tide framework for application to river level data. We note that this would not remove the need for a secondary port in NS_Tide which we regard as a significant innovation of our approach.

Figure 3, Panel (A) shows the relative location of each of these data sources. Figure 3, Panel (B) shows the observed sea level below the Chester Weir, alongside concurrent tide gauge data from Liverpool, and measured water levels of the River Dee upstream of the weir. While there is a relationship between the upstream river level, tidal signal, and the water elevation below the Weir, the relationship is not linear and crucially, non-harmonic (Matte et al., 2014).

To test the performance of both methods, RTide, and IRLS HA are trained using 10 months of reference data (≈20,000 observations) from beneath the Chester Weir. Here, IRLS HA is utilized as the data is strongly non-stationary. Predictions are then computed for the subsequent 2 months (November to December). The 10 months hindcasts and 2 months forecasts are shown in Figure 3, Panel (C). In addition to the spherical harmonic input functions, RTide is provided with the real-time upstream river levels. Critically, unlike other nonstationary methods, RTide does not require the inclusion of an adjacent gauge without fluvial effects.

Visually it can be seen that the RTide prediction is better able to track the non-stationary fluctuations induced by river outflow than HA. This is confirmed by the coefficients of determination $\left(R^2\right)$ of 0.922 m and −.321 for RTide and HA respectively on the test data set. This result is not surprising as the tidal dynamics are heavily modulated by the tidal river as will be shown in the subsequent SHAP analysis. By removing the assumption of stationary behavior, RTide can easily learn the coupled response of sea-levels to non-stationary forcing. We note that the use of IRLS HA in this case study is meant to illustrate the severe biasing of harmonic predictions when applied to areas exhibiting nonstationarity and is not meant as a direct benchmark.

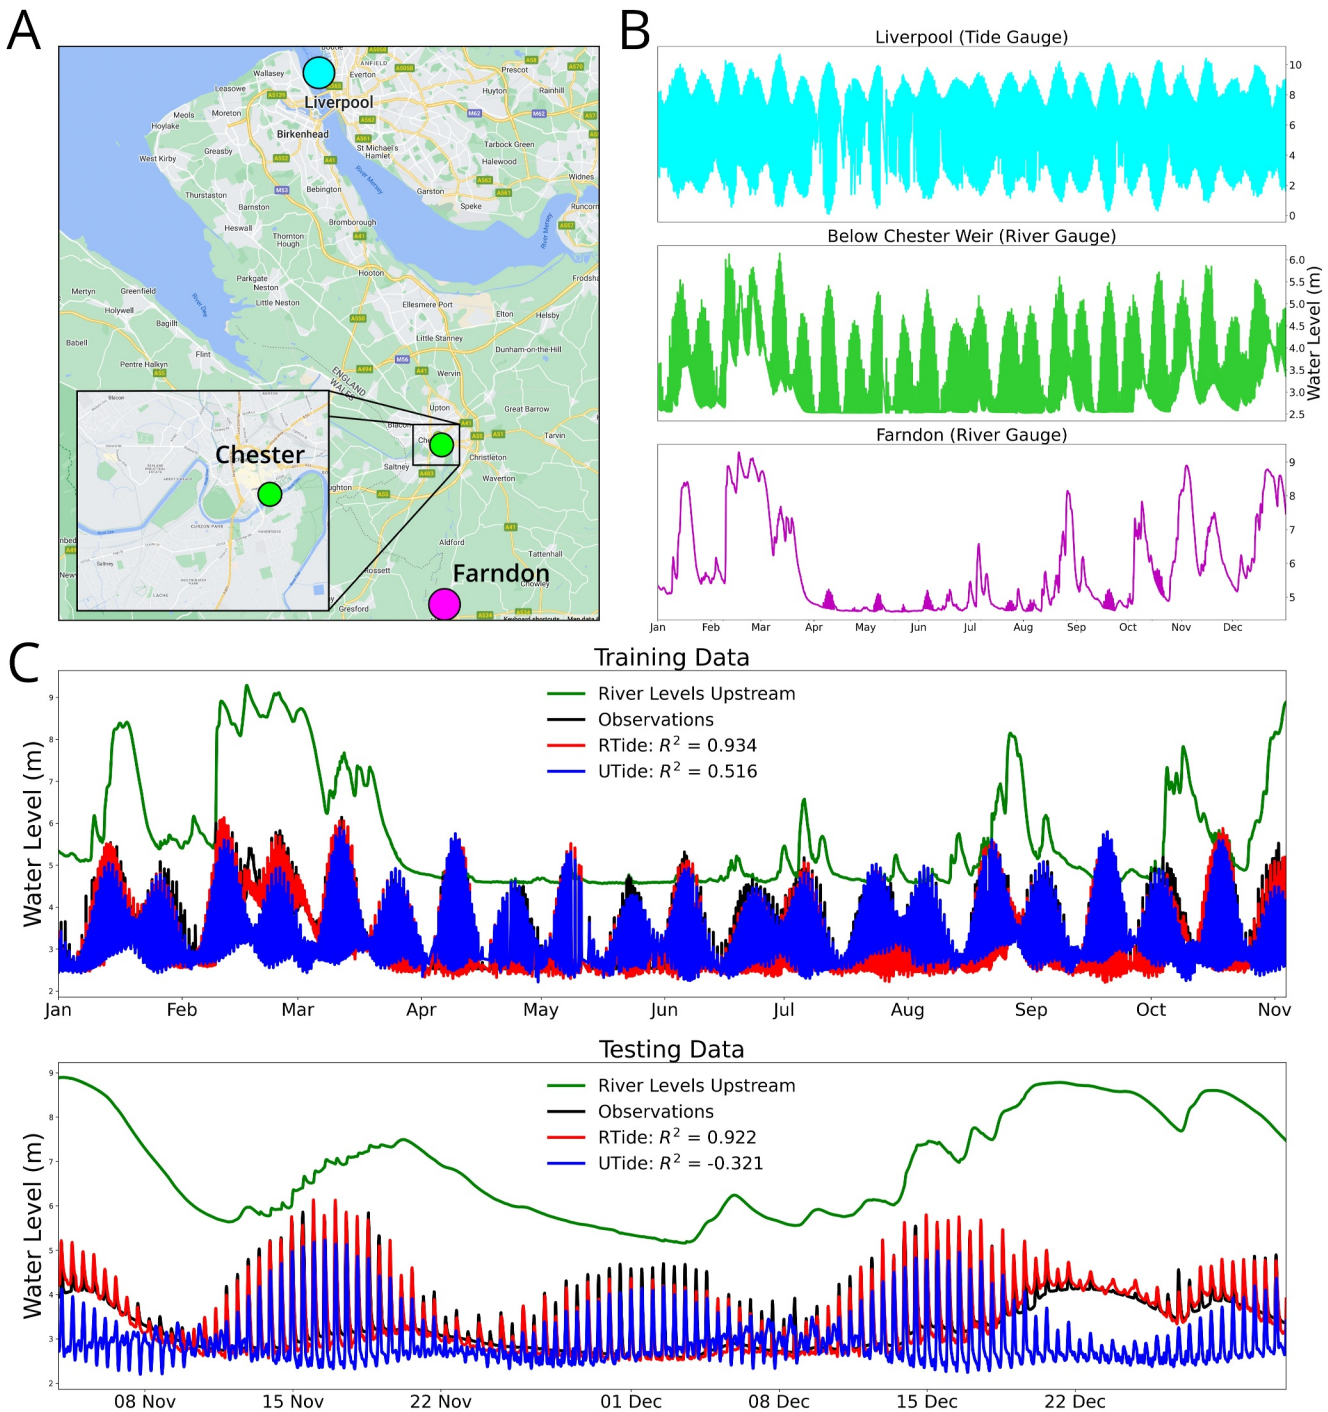

Figure 3. Comparisons between RTide and IRLS HA predictions of tidal river levels beneath Chester Weir, UK. Panel (a) shows the relative location of each data source. Panel (b) shows the time series of each data source. Panel(s) (c) show RTide and UTide predictions over the full year with the top panel being hindcasts (Training Data), and the bottom panel showing forecasts (Testing Data). The coefficient of determination, R^2, of each models predictions is given. The only inputs into RTide are the upstream river levels, and gravitational input functions.

RTide demonstrates balanced performance across the training and test sets with R^2 values of 0.934 and 0.922 respectively (the difference is not significant based on t-test with $\alpha = 0.05$). This indicates that RTide has learned the localized and nonlinear response of the sea-level beneath the Chester weir to the combined tidal and fluvial forcing. The learned model can therefore be used to inform decisions and quantify the risk associated with

different water management scenarios. This is important as the gauge beneath the Chester weir has since been decommissioned. Unlike autoregressive machine learning approaches (Qin et al., 2023; Xie et al., 2023), our approach enables the accurate and continued prediction of water-levels in this region based on this historical data.

In addition to the absolute improvements in predictive accuracy, the model can also provide insights into the underlying dynamics (e.g., what aspects of the tidal signal interact with fluvial effects? How does that impact tidal propagation?).

This can be seen in Figure 4, and Figure S3 in Supporting Information S1. Here the RTide model is analyzed using the SHAP analysis procedure outlined above. Dependence plots are shown which visualize the learned response *distributions* between the individual input functions and the associated sea-level response beneath the Chester Weir. These distributions are produced by interactions between input functions. Indeed, if no such interactions exist then an analytic form could be approximated, but care should be taken.

Each plot illustrates the individual contribution of each physical driver, and the nonlinear interaction between them. As noted, the spread in the derived response relationships is a consequence of interactions with other inputs, which when plotted together, can shed light on how these physical drivers interact. For example, Figure 3 shows the learned response to the upstream river levels, and exhibits a near cubic relationship $\left(R^2 = 0.987\right)$ with relatively little spread in the associated distribution reflected by the ratio of the standard deviation of the fit model's residuals σ_{fit} compared to the standard deviation of the data σ_{data} ($\sigma_{\text{fit}}/\sigma_{\text{data}} = 0.11$). Inspection of the distribution of the color-mapped astronomical tide indicates that at low upstream river levels ($\leq$7 m) interactions with the tide create a larger sea-level response. In contrast, for high upstream river levels ($\geq$7 m), and thus high river outflow, the incoming tide appears to have a damping effect. These learned responses nicely align with theory (Jay & Kukulka, 2003).

Figure 4 illustrates the impact the river outflow has on the astronomical tide by comparing the RTide derived gravitational response between Liverpool and Chester for individual input functions and time-lags. Gravitational$_n^m$ refers to the m^{th} order and n^{th} degree gravitational spherical harmonic input function. The Gravitational$_2^2$ spherical harmonic shown in Figure 4 corresponds to the semi-diurnal species and thus primarily drives its response. The tides at Liverpool have some fluvial influence due to the river Mersey but sit adjacent to the open ocean and can thus be used to evaluate the impact of the introduction of river outflow. It can be seen that the river outflow modulates the tidal response in several ways. First, even for time-lags associated with the semi-diurnal tidal response, we see the upstream river levels induce a spread of the sea-level response for the Chester gravitational response distributions, as opposed to the relatively minimal spread of the Liverpool responses. This spread exhibits a clear gradient when considered with the river outflow. Here, high river levels upstream appear to damp the gravitational response leading to larger sea-level responses at lower river levels and vice versa.

The second factor affecting the response distributions is nonlinear distortion. The Liverpool responses can be seen to be near linear based on linear regressions with $R^2 > 0.99$. In contrast, the Chester response distributions are distorted and non-linear (regression R^2 of 0.62, 0.93, and 0.16). Figure S3 in Supporting Information S1 shows the effect of these modulations on the overall gravitational input functions separated by different ranges of upstream river level. These varied nonlinear relationships highlight the difficulty encountered in the classic approach – where one must prescribe these relationships a priori. Future work will look to explore these relationships further, and to connect the observed response distribution modulation to the extensive work which has been done for HA (Jay & Kukulka, 2003; Matte et al., 2014).

While a river-level forecasting model is needed to operationalize these tools, considerable work in the field of hydrology has looked to solve this problem (Galavi et al., 2013; See & Openshaw, 1999, 2000). Furthermore, while results could be improved by using river outflow as an additional forcing input, the proposed approach has the advantage of being applicable to remote-sensing data. This will be useful for the development of global estuarine and tidal river models on the upcoming NASA Surface Water Ocean Topography (SWOT) mission. Further study is needed using data where both river outflow measurements and uncontaminated secondary ports are available to directly compare the predictive capabilities of RTide with NS_TIDE.

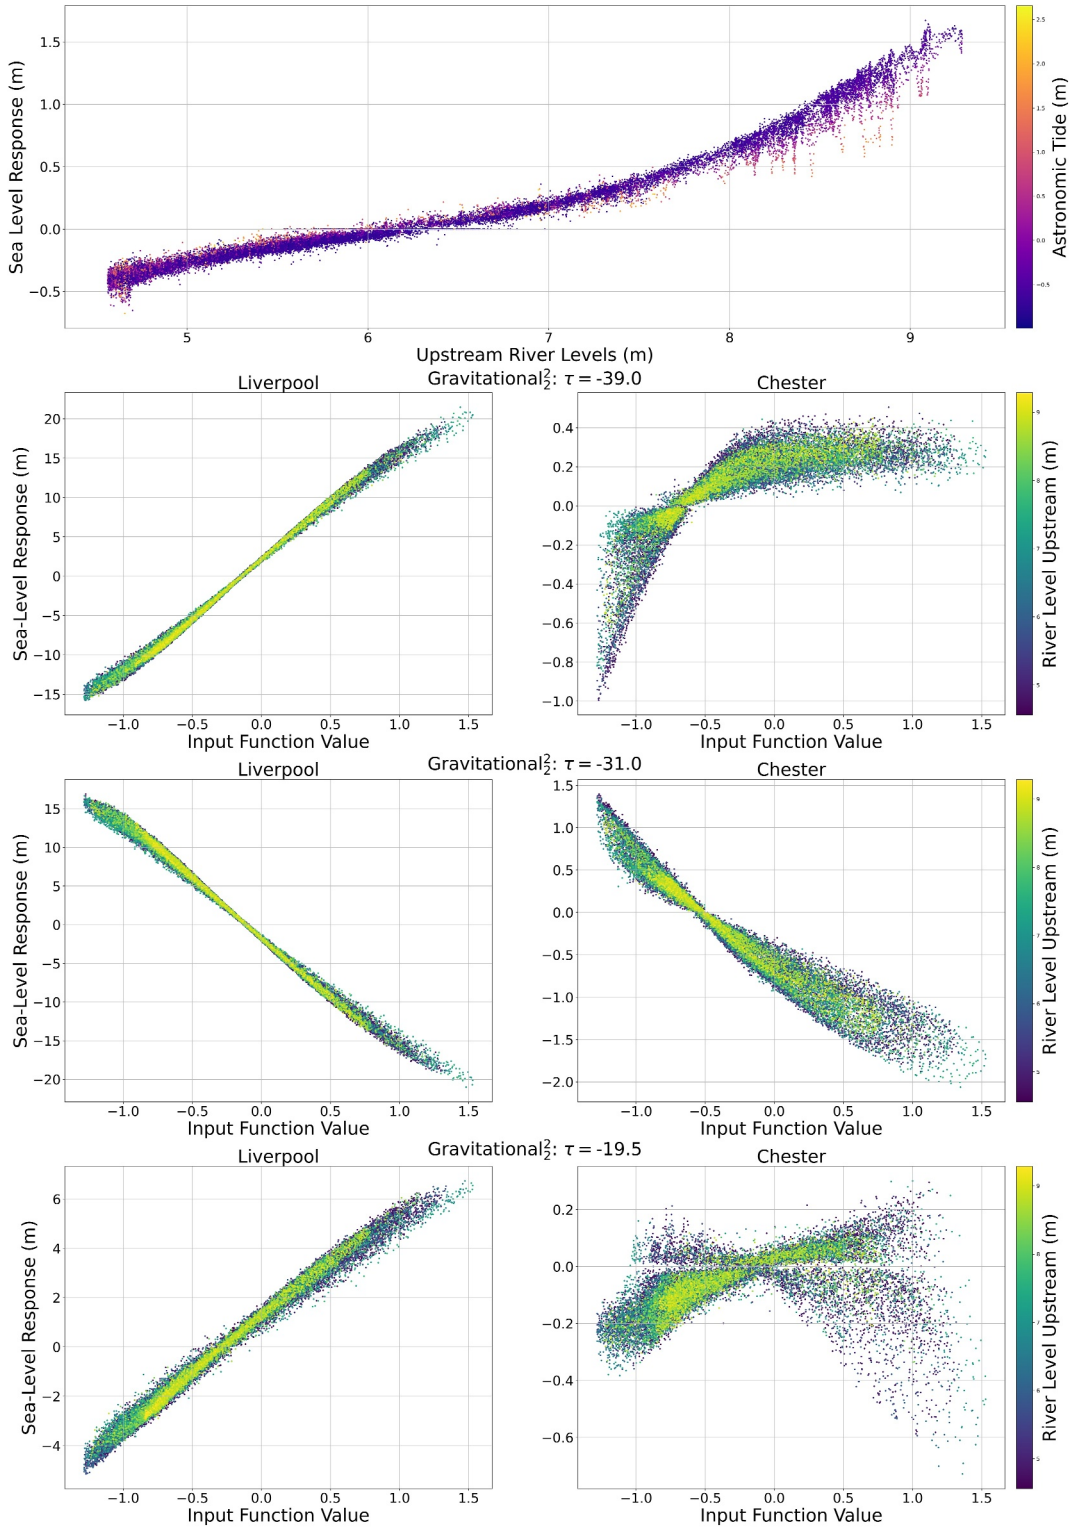

Figure 4. SHAP derived response distributions of the Chester, UK sea-level. (Top Panel) Learned response distribution to upstream river levels. Data points are colored according to the model predicted astronomical tide. (Bottom Panels) Comparisons of learned response distributions to the second order and degree Gravitational input functions at associated time-lags τ between Liverpool, UK and Chester, UK. Chester data is colored according to the upstream river levels at the time of measurement. Corresponding input function, spherical harmonic, and time-lag are given above each subplot. Associated values for all plots are computed using SHAP for the entire year-long time series.

4.2. Meteorological Forcing

As sea levels continue to rise, coastal flooding and storm surges will become more frequent (Tebaldi et al., 2012). While major flooding events, primarily driven by hurricanes and typhoons, are of great importance, the incidence of minor flooding will proliferate in the coming years. Often referred to as nuisance flooding, these events rarely result in loss of life but place severe strain on public infrastructure and property (Moftakhari et al., 2018). In fact, the cumulative cost of these frequent events can exceed that of less frequent, major flooding events (Moftakhari et al., 2017). The drivers of nuisance floods are generally more subtle and are a result of the confluence of astronomical tides, wind, and weather.

Classically, the task of surge prediction is carried out using numerical methods (Kohno et al., 2018). These hydrodynamic models are computationally intensive and suffer from over-simplistic assumptions that prevent the consideration of coupled processes (Jelesnianski, 1992; Luettich et al., 1992). In particular, many operational models treat tides and surges as separable due to computational limitations. Several studies have shown that treating surge and tide as separate can lead to erroneous predictions (Gou et al., 2023; Horsburgh & Wilson, 2007; Rego & Li, 2010). Additionally, the high computational cost of numerical models is impractical for ensemble forecasting and regions with poor data quality or access to computational resources. To combat these limitations we propose the ML Response Framework as a lightweight tool for learning the localized response to meteorological forcing. We stress that the present study is not developing an alternative to the operational surge models, but rather illustrating the potential of our approach for further applications in surge modeling and understanding the role of meteorological forcing.

We test the proposed method using data from Money Point, Virginia, USA over a 5-year period from 1 January 2014, to 31 December 2018. Money Point is situated within the Chesapeake Bay's inner reaches, along the Elizabeth River's banks. Due to the region's low elevation and proximity to other waterways, flooding is a common occurrence. The tide gauge at Money Point (ID:8639348) is paired with a small meteorological station measuring the ambient temperature, barometric pressure, and wind speed and direction. While there are many other forcing variables that could be included, we focus on these variables as they are the primary drivers of surge events and are generally available from meteorological stations and local weather models. RTide and HA models are trained on 365 days of reference data from 2014 to 2015. To compare the importance of different forcing inputs, Table 1 presents performance indicators for the subsequent 4-year predictions using different combinations of input variables. RTide hyperparameters were kept constant between runs, meaning that additional performance could likely be achieved through optimization. It can be seen that the predictive accuracy of RTide improves with the inclusion of additional meteorological forcing. The results indicate that the inclusion of wind has the single largest impact on predictive accuracy at this location. Interestingly, we found that the inclusion of gust information slightly reduces long term performance, however, it improved the prediction of some extreme surge events. The reduced long term performance is likely a consequence of being highly correlated with the wind variables. While we demonstrate significant improvements over the (Gravitational + Radiational) RTide and HA methods, we expect further improvements could be achieved by incorporating additional forcing variables and spatial information. The inclusion of spatial forcing is critical for predicting surges originating externally such as those driven by hurricanes. Furthermore, the present analysis has not looked to optimize the lags for this task as we wish to test the "off-the-shelf" performance of RTide. This represents an important area of future work and can likely improve performance significantly.

Evaluation of the SHAP values for the model's prediction of a flooding event can help to explain the event's physical drivers. We illustrate this using a 5-day surge event from May 4th to 9th, 2016, characterized by sustained on-shore winds. The minor flood level threshold is designated by NOAA's Weather Forecasting Office and is empirically defined by occurrences of local flooding. For Money Point, VA this is 54 cm above the mean higher high water. Figure 5, Panel (A) explains the influence of each forcing input on the RTide prediction of the first minor flooding event when trained on pressure, wind, temperature, and wind gusts. It can be seen that the combination of high astronomical forcing and sustained Southern winds over the previous 72 hr contributed to this event. This information can be useful to local environmental agencies as the tide-surge response to meteorological forcing is highly location and time-dependent.

Panel (B) compares RTide and HA predictions for the full 5-day surge event. Visually, the RTide prediction shows strong agreement with the observed sea level, particularly during the peak flood events. While predicting these occurrences is clearly of the highest importance, prediction accuracy decreases during low-tide events. This

Table 1
Comparisons of HA and RTide Performance From 1 January 2015, to 31 December 2018 in Money Point, VA (NOAA Station ID: 8639348)

Model	RMSE	MAE	Max error	R^2
HA	0.19097	0.14342	1.23389	0.6528
Tidal	0.1927	0.14272	1.28285	0.65712
Tidal + Pressure	0.17935	0.13565	1.08744	0.69942*
Tidal + Wind	0.17811	0.1375	**0.8686**	0.72835*
Tidal + Temp	0.20222	0.15484	1.19162	0.64787
Tidal + Wind + Pressure	0.16304	0.12591	0.90245	0.7852*
Tidal + Wind + Pressure + Temp	**0.1364**	**0.10443**	0.87945	**0.84161***
Tidal + Wind + Pressure + Temp + Gust	0.14062	0.10586	1.12529	0.83149*

Note. Here "Tidal" refers to RTide using just gravitational and radiational input functions. Additional models are named according to the additional meteorological variables they are trained on. Best values are provided in bold. Statistically significant improvements in R^2 over HA are indicated using * and are computed via t-test with $\alpha = 0.05$.

is likely a consequence of only considering localized meteorological forcing. As was demonstrated in Section 4.1, we anticipate the inclusion of additional spatial information will further improve predictive accuracy. Experiments with the inclusion of positive meteorological lags showed promise in predicting external surges (e.g., surges originating from elsewhere on the shelf). However, further analysis is needed to confirm these findings.

Panel (C) shows the evolution of the four meteorological variables over the 5-day surge event horizon. These reveal a key challenge in predicting nuisance floods; flooding results from the coincidence of high tide events with both the chronic build-up and acute effects of meteorological forcing. The acute impacts of meteorological forcing can be seen in Panel (D), which shows distributions of the model's learned dependence of the sea-level to real-time meteorological forcing taken over a 6 month period. Analysis of these distributions shows that the learned RTide feature importance is consistent with our physical intuition. There is a clear, and significant, inverse linear relationship between barometric pressure and the localized sea-level response. Linear regression of the learned sea-level response to pressure yields a slope of -0.83 cm/mbar which is in good agreement with both theoretical and empirical studies (Fu & Pihos, 1994). Some discrepancy may arise from the fact that we only consider the barometric pressure at the gauge, and not the average over the region. Additionally, we see that elevated wind gusts in the North-East direction have the most significant impact on real-time sea-level with Easterly gusts. This finding makes sense as the station sits on the South-Western point of a jetty. It is interesting that the wind exhibits slight negative sea-level responses (-0.6 cm/kmh^{-1}), and the gusts large positive responses (2 cm/kmh^{-1}). However, a simple explanation is that the total response to wind is the sum of the response from both sustained winds and gusts. Adding these two components shows a total wind response which is approximately linear with a response slope of (1.14 cm/kmh^{-1}) for the U component (Figure S4 in Supporting Information S1). The spread in the distribution (nonlinearity) seems to increase with the difference between sustained wind and gust as shown by the colormap. The model response shows ambient temperature has essentially no impact (-0.05 cm/C°). This illustrates the value of SHAP analysis in explaining models with multiple partially correlated inputs which proved difficult to define or interpret in Cartwright's original surge analysis (Cartwright, 1968). Through our testing, we found these relationships became more clearly defined when models were trained on longer reference series and regularized. Future work will look to verify the nature of these dependencies on longer reference series and to compare them to different stations.

In this application, we use the observed meteorological variables as direct inputs into our model. For an operational model, regional forecasts would need to be substituted. This "lightweight" surge-prediction framework shows promise as an alternative to conventional numerical methods for communities lacking computational resources and in predicting nuisance flooding. Notably, none of the examples presented here utilize the observed water-levels as input in contrast with the machine learning approaches previously employed for this task (Qin et al., 2023). Inclusion of recent or real-time water levels as exogenous inputs, termed "self-prediction" by Cartwright (1968), can significantly improve performance. However, by avoiding this requirement, our approach

10.1029/2024JH000525

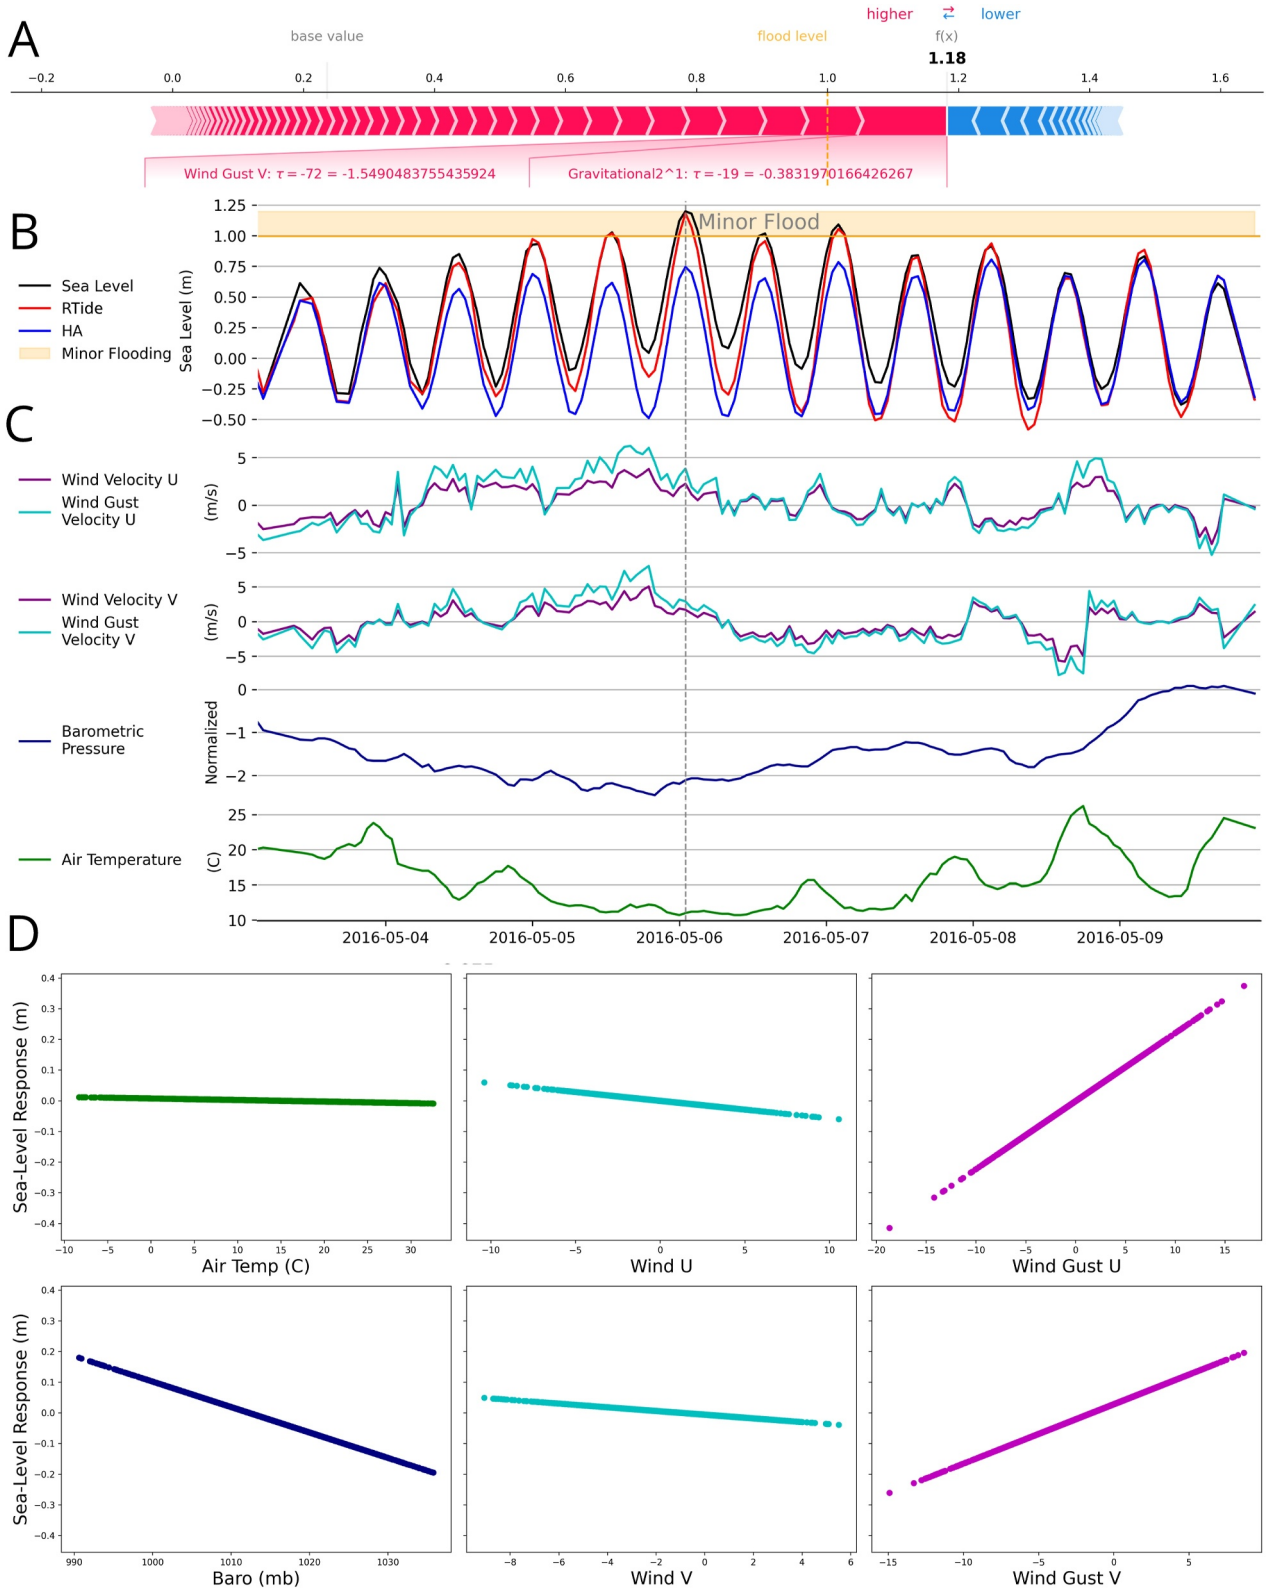

Figure 5. 5-day nuisance flooding event in Money Point, VA, USA. Panel (a) shows the explained surge prediction at the time instance labeled "Minor Flood." Panel (b) shows RTide and HA predictions of the observed sea level. Panel (c) shows the subset of meteorological variables associated with the 5-day flooding event and RTide predictions. Panel (d) shows distributions of SHAP sea-level dependence by real-time meteorological variables. Distributions are shown for 4200 predictions over 180 days from January through June of 2016. RTide uses all meteorological variables shown.

is readily applied to ungauged locations and does not suffer from potential observational issues (Qin et al., 2023). Additionally, while such an approach naturally falls out of the response framework, it is not possible in a harmonic context.

Regardless, this case study illustrates that sea-level predictions and subsequent analysis at any tide gauge station with access to meteorological data can be improved by RTide. This is promising for applications in developing regions, where limited data and computational resources needed to run conventional large numerical models exist. As MSL rise continues, the margin of error in surge prediction will decrease. Hence, improving our understanding of the localized influences of meteorological forcing will be critical to adaptation efforts moving forward.

5. Discussion

The new methods of analysis introduced by automating the response method have applications in oceanography, coastal engineering, marine renewable energy, and the study of anthropogenic climate change. In this work, we demonstrate several applications of our approach to contemporary problems in tidal analysis and prediction. While we have rigorously compared the performance of RTide against HA for these use cases, it remains to be said how our method fits in with conventional tidal analysis. In this section, we discuss the strengths and limitations of our method and propose several avenues of future work with these in mind.

While our framework demonstrates excellent performance for conventional tidal prediction problems, we stress that for the majority of these conventional use cases, standard HA remains the preferred choice. This is a consequence of the myriad of benefits garnered by HA: compactness, interpretability, reliability, computation time, etc. While the physics-informed construction of our method and development of new techniques for analysis have looked to mitigate the inherent opaqueness and stochasticity of ML methods, these limitations must be acknowledged.

For "messy" applications in more complex tidal regimes, the flexibility and non-parametric formulation of our method become advantageous. As illustrated in Section 3.1, conventional HA can cause problems when applied to signals contaminated with non-astronomical forcing. One such example is the biasing of harmonic coefficients from radiational effects. Known as the radiational double-counting problem, this can lead to significant errors for operational surge forecasts (Williams, Irazoqui Apecechea, et al., 2018; Williams, Saulter, et al., 2018). While non-stationary extensions to HA have been developed and applied for the analysis and prediction of sea-levels in tidal rivers and estuaries (George & Simon, 1984; Matte et al., 2014), these approaches encounter several challenges. First, both of these approaches require an adjacent "uncontaminated" reference series. As noted above, for many regions globally, particularly in the Global South, this is not feasible. Additionally, with regard to NS_Tide one must first understand the physics of the interaction between tide and external forcing, which, for many important problems, is not yet well defined (storm surges, sea level rise, intertidal regions, etc.).

The non-parametric nature of the ML approach removes this requirement and enables the study of these phenomena without imposing any assumptions regarding separability, stationarity, or the structure of interaction kernels. The flexibility in input forcing selection garnered by these characteristics makes the testing of multivariate forcing simple and efficient. Future work will look to combine these capabilities with symbolic regression in order to derive analytic approximations of the learned response models. The insights could then be infused into conventional non-stationary HA to leverage the aforementioned benefits of a simple regression model. We remark that an analogous framework could be conceived of for HA if a non-parametric term could be simultaneously estimated (e.g., NS_Tide without predefined basis function), similar to the use of "universal differential equations" (Rackauckas et al., 2020). As shown in Section 3.1 we believe that simultaneous estimation is critical to recovering the true physical processes.

In addition to the results shown in Sections 4.1 and 4.2, the approach to analyzing the learned model using SHAP provides many further avenues with which to study the impacts of non-gravitational forcing. Indeed, Figures 4 and S2 in Supporting Information S1 provide a first look at how this can be used to understand how non-gravitational forcing modulates the astronomical tide and vice versa. Additionally, Section 4.2 demonstrated how the SHAP values can be used to recover known physical processes, such as the inverse barometer effect, even in the presence of multiple drivers. This constitutes an exciting new area of study as countless phenomena interact and modulate the astronomical tide but are difficult to study and predict using conventional methods. For example, interactions of tides with mean sea-level (and thus implications for sea level rise)

(Cartwright, 1968), and even the impact of barrages, are both highly salient for coastal resilience efforts. However, current efforts to include these interactions in engineering practice are generally limited to statistical estimation or computationally intensive numerical simulation (Idier et al., 2019). Initial testing has shown our approach to be promising for both such applications and will be pursued in the future. Given their physical interpretability, we believe SHAP values can help enable new explorations into the dynamics of the ocean's response to various forcing, as demonstrated.

The present study has primarily conducted our analysis of the learned RTide models within the time-domain. As shown in (Cartwright, 1968; Munk & Cartwright, 1966), the frequency interpretation of response models is similarly rich. In fact, Munk and Cartwright found the admittance function $Z(f)$ of the learned response weights to be more meaningful than the response weights themselves. It is possible to obtain the equivalent Volterra series directly from the trained networks (Wray & Green, 1994). This approach could be useful in (a) studying the nature of the learned interactions and (b) translating these models back into a familiar response regression. The learned admittance $Z(f)$ can then be computed and analyzed in the standard way. Alternatively, as described in Section 2.3, SHAP values can serve as a transformation from our learned ML model to a generalized linear model of binary variables. Hence, future work will look to exploit both of these approaches to analyze the approximate admittance of learned RTide models and compare them with the original method.

The accurate prediction of estuarine tides remains a significant challenge due to distortion of the tidal signal which can give rise to asymmetric and non-harmonic profiles (sawtooth, bore formation, etc.). Inaccuracies from these predictions constitute one of the largest sources of error for the UK operational surge model (Williams, Saulter, et al., 2018). As demonstrated, a response approach can readily capture and predict these modulations. Future work will look to operationalize these forecasts for integration into these tools.

6. Conclusion

The machine learning approach advanced in this manuscript automates the tidal response method and makes its application to nonstationary tidal processes simple. RTide is thus the first general purpose analysis and prediction tool for arbitrary tidal processes.

By utilizing a class of neural networks which are equivalent to the Volterra series, the approach retains equivalence to the classical response method, and under certain conditions, the analytic form can be recovered exactly (Marmarelis & Zhao, 1997). Constraining these networks within the impulse-response framework:

1. Removes the need to define interactions between input functions.
2. Allows the network to learn the higher order impulse-response relationships.
3. Enables prediction of non-stationary sea-levels without requiring an adjacent reference station removed from such influences.
4. Makes easy the inclusion, quantification, and analysis of interactions between non-astronomical forces with tides.

An implementation of our approach is provided in the easy-to-use RTide Python package which produced all of the provided case-studies in just three lines of code (Monahan, 2024). Following the tradition set by Munk and Cartwright, we conclude with the remarks of Hilaire Belloc: “When they pontificate on the tides it does no great harm, for the sailorman cares nothing for their theories, but goes by real knowledge.” It is our hope that RTide can help to explain this “real knowledge.”

Data Availability Statement

Code and data to replicate all applications in this manuscript is given in the Zenodo repository (Monahan, 2025). Jupyter notebooks containing the model implementation and subsequent analysis are included for all figures. NOAA tide gauge data (Section 4.2) can also be accessed freely via NOAA Tides and Currents repository tidesandcurrents.noaa.gov. The tide and river gauge data in Section 4.1 is provided courtesy of the UK National Oceanography Centre.

Acknowledgments
TM and TT acknowledge support from the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a Schmidt Futures program. We would like to thank Phil Woodworth for the insightful discussions and feedback on our manuscript. Additionally, we would like to thank Jeff Polton, and the UK National Oceanography Centre for providing data, guidance, and helpful conversations.

References

Andersson, C., Ribeiro, A. H., Tiels, K., Wahlström, N., & Schön, T. B. (2019). Deep convolutional networks in system identification. In *2019 IEEE 58th conference on decision and control (CDC)* (pp. 3670–3676).

Brunton, S. L., Noack, B. R., & Koumoutsakos, P. (2020). Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, *52*(1), 477–508. https://doi.org/10.1146/annurev-fluid-010719-060214

Cartwright, D. E. (1968). A unified analysis of tides and surges round north and East Britain. *Philosophical Transactions of the Royal Society of London - Series A: Mathematical and Physical Sciences*, *263*(1134), 1–55.

Cartwright, D. E., & Ray, R. (1990). Oceanic tides from Geosat altimetry. *Journal of Geophysical Research: Oceans*, *95*(C3), 3069–3090. https://doi.org/10.1029/jc095ic03p03069

Chen, K., Kuang, C., Wang, L., Chen, K., Han, X., & Fan, J. (2021). Storm surge prediction based on long short-term memory neural network in the East China Sea. *Applied Sciences*, *12*(1), 181. https://doi.org/10.3390/app12010181

Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhye, H., et al. (2016). Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems* (pp. 7–10).

Codiga, D. L. (2011). Unified tidal analysis and prediction using the UTide Matlab functions.

De Oliveira, M. M., Ebecken, N. F. F., De Oliveira, J. L. F., & de Azevedo Santos, I. (2009). Neural network model to predict a storm surge. *Journal of Applied Meteorology and Climatology*, *48*(1), 143–155. https://doi.org/10.1175/2008jamc1907.1

Donini, J., Hollebone, B., & Lever, A. (1977). The derivation and application of normalized spherical harmonic Hamiltonians. *Progress in Inorganic Chemistry*, *22*, 225–261.

Doodson, A. T. (1921). The harmonic development of the tide-generating potential. *Proceedings of the Royal Society of London - Series A: Containing Papers of a Mathematical and Physical Character*, *100*(704), 305–329.

Doodson, A. T. (1929). *Report on Thames floods (No. 7)*. HM Stationery Office.

Doodson, A. T. (1958). Oceanic tides. In *Advances in geophysics* (Vol. 5, pp. 117–152). Elsevier. https://doi.org/10.1016/s0065-2687(08)60077-1

Egbert, G. D., & Ray, R. D. (2017). Tidal prediction. *Journal of Marine Research*, *75*(3), 189–237. https://doi.org/10.1357/002224017821836761

Fowler, A., & Kember, G. (1993). Delay recognition in chaotic time series. *Physics Letters A*, *175*(6), 402–408. https://doi.org/10.1016/0375-9601(93)90991-8

Fu, L.-L., & Pihos, G. (1994). Determining the response of sea level to atmospheric pressure forcing using topex/poseidon data. *Journal of Geophysical Research: Oceans*, *99*(C12), 24633–24642. https://doi.org/10.1029/94jc01647

Galavi, H., Mirzaei, M., Shul, L. T., & Valizadeh, N. (2013). Klang river–level forecasting using ARIMA and ANFIS models. *Journal of the American Water Works Association*, *105*(9), E496–E506. https://doi.org/10.5942/jawwa.2013.105.0106

Gan, M., Chen, Y., Pan, S., Li, J., & Zhou, Z. (2019). A modified nonstationary tidal harmonic analysis model for the Yangtze estuarine tides. *Journal of Atmospheric and Oceanic Technology*, *36*(4), 513–525. https://doi.org/10.1175/jtech-d-18-0199.1

George, K., & Simon, B. (1984). The species concordance method of tide prediction in estuaries. *International Hydrographic Review*.

Gou, X., Liang, H., Cai, T., Wang, X., Chen, Y., & Xia, X. (2023). The impact of coastline and bathymetry changes on the storm tides in zhejiang coasts. *Journal of Marine Science and Engineering*, *11*(9), 1832. https://doi.org/10.3390/jmse11091832

Groves, G. W., & Reynolds, R. W. (1975). An orthogonalized convolution method of tide prediction. *Journal of Geophysical Research*, *80*(30), 4131–4138. https://doi.org/10.1029/jc080i030p04131

Hoitink, A., & Jay, D. A. (2016). Tidal river dynamics: Implications for deltas. *Reviews of Geophysics*, *54*(1), 240–272. https://doi.org/10.1002/2015rg000507

Horsburgh, K., & Wilson, C. (2007). Tide-surge interaction and its role in the distribution of surge residuals in the North Sea. *Journal of Geophysical Research: Oceans*, *112*(C8). https://doi.org/10.1029/2006jc004033

Idier, D., Bertin, X., Thompson, P., & Pickering, M. D. (2019). Interactions between mean sea level, tide, surge, waves and flooding: Mechanisms and contributions to sea level variations at the coast. *Surveys in Geophysics*, *40*(6), 1603–1630. https://doi.org/10.1007/s10712-019-09549-5

Idier, D., Dumas, F., & Muller, H. (2012). Tide-surge interaction in the English channel. *Natural Hazards and Earth System Sciences*, *12*(12), 3709–3718. https://doi.org/10.5194/nhess-12-3709-2012

Jay, D. A., & Kukulka, T. (2003). Revising the paradigm of tidal analysis–the uses of non-stationary data. *Ocean Dynamics*, *53*(3), 110–125. https://doi.org/10.1007/s10236-003-0042-y

Jelesnianski, C. P. (1992). *SLOSH: Sea, lake, and overland surges from hurricanes* (Vol. 48). US Department of Commerce, National Oceanic and Atmospheric Administration.

Judd, K., & Mees, A. (1998). Embedding as a modeling problem. *Physica D: Nonlinear Phenomena*, *120*(3–4), 273–286. https://doi.org/10.1016/s0167-2789(98)00089-x

Keers, J. (1968). An empirical investigation of interaction between storm surge and astronomical tide on the east coast of Great Britain. *Deutsche Hydrographische Zeitschrift*, *21*(3), 118–125. https://doi.org/10.1007/bf02235726

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kohno, N., Dube, S. K., Entel, M., Fakhruddin, S., Greenslade, D., Leroux, M.-D., et al. (2018). Recent progress in storm surge forecasting. *Tropical Cyclone Research and Review*, *7*(2), 128–139.

Kratzert, F., Klotz, D., Shalev, G., Klambauer, G., Hochreiter, S., & Nearing, G. (2019). Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets. *Hydrology and Earth System Sciences*, *23*(12), 5089–5110. https://doi.org/10.5194/hess-23-5089-2019

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., et al. (2023). Learning skillful medium-range global weather forecasting. *Science*, *382*(6677), 1416–1421. https://doi.org/10.1126/science.adi2336

Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., & Talwalkar, A. (2017). Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, *18*(1), 6765–6816.

Lobo, M., Jay, D. A., Innocenti, S., Talke, S. A., Dykstra, S. L., & Matte, P. (2024). Implementing superresolution of nonstationary tides with wavelets: An introduction to CWT_Multi. *Journal of Atmospheric and Oceanic Technology*, *41*(10), 969–989. https://doi.org/10.1175/jtech-d-23-0144.1

Luettich, R. A., Westerink, J. J., & Scheffner, N. W. (1992). *ADCIRC: An advanced three-dimensional circulation model for shelves, coasts, and estuaries. Report 1, theory and methodology of ADCIRC-2DD1 and ADCIRC-3DL*. Coastal Engineering Research Center (US).

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, *30*.

Lundberg, S. M., Nair, B., Vavilala, M. S., Horibe, M., Eisses, M. J., Adams, T., et al. (2018). Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature Biomedical Engineering*, *2*(10), 749–760. https://doi.org/10.1038/s41551-018-0304-0

Marmarelis, V. Z., & Zhao, X. (1997). Volterra models and three-layer perceptrons. *IEEE Transactions on Neural Networks*, *8*(6), 1421–1433. https://doi.org/10.1109/72.641465

Matte, P., Secretan, Y., & Morin, J. (2014). Temporal and spatial variability of tidal-fluvial dynamics in the St. Lawrence fluvial estuary: An application of nonstationary tidal harmonic analysis. *Journal of Geophysical Research: Oceans*, *119*(9), 5724–5744. https://doi.org/10.1002/2014jc009791

Moftakhari, H. R., AghaKouchak, A., Sanders, B. F., Allaire, M., & Matthew, R. A. (2018). What is nuisance flooding? Defining and monitoring an emerging challenge. *Water Resources Research*, *54*(7), 4218–4227. https://doi.org/10.1029/2018wr022828

Moftakhari, H. R., AghaKouchak, A., Sanders, B. F., & Matthew, R. A. (2017). Cumulative hazard: The case of nuisance flooding. *Earth's Future*, *5*(2), 214–223. https://doi.org/10.1002/2016ef000494

Monahan, T. (2024). RTide. Retrieved from https://github.com/thomasmonahan/RTide.GitHub

Monahan, T. (2025). Replication materials for rtide: Automating the tidal response method [Dataset]. *Zenodo*. https://doi.org/10.5281/zenodo.15295189

Monahan, T., Tang, T., & Adcock, T. A. A. (2023). A hybrid model for online short-term tidal energy forecasting. *Applied Ocean Research*, *137*, 103596. https://doi.org/10.1016/j.apor.2023.103596

Monahan, T., Tang, T., Roberts, S., & Adcock, T. A. (2025). Tidal corrections from and for SWOT using a spatially coherent variational Bayesian harmonic analysis. *Journal of Geophysical Research: Oceans*, *130*(3), e2024JC021533. https://doi.org/10.1029/2024jc021533

Munk, W. H., & Cartwright, D. E. (1966). Tidal spectroscopy and prediction. *Philosophical Transactions of the Royal Society of London - Series A: Mathematical and Physical Sciences*, *259*(1105), 533–581.

Pan, H., Lv, X., Wang, Y., Matte, P., Chen, H., & Jin, G. (2018). Exploration of tidal-fluvial interaction in the Columbia River estuary using S_TIDE. *Journal of Geophysical Research: Oceans*, *123*(9), 6598–6619. https://doi.org/10.1029/2018jc014146

Parker, B. B. (2007). *Tidal analysis and prediction*. NOAA, NOS Center for Operational Oceanographic Products and Services.

Prandle, D., & Wolf, J. (1978). The interaction of surge and tide in the North Sea and River Thames. *Geophysical Journal International*, *55*(1), 203–216. https://doi.org/10.1111/j.1365-246x.1978.tb04758.x

Qin, Y., Su, C., Chu, D., Zhang, J., & Song, J. (2023). A review of application of machine learning in storm surge problems. *Journal of Marine Science and Engineering*, *11*(9), 1729. https://doi.org/10.3390/jmse11091729

Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., et al. (2020). Universal differential equations for scientific machine learning. arXiv preprint arXiv:2001.04385.

Rala Cordeiro, J., Raimundo, A., Postolache, O., & Sebastião, P. (2021). Neural architecture search for 1D CNNS—Different approaches tests and measurements. *Sensors*, *21*(23), 7990. https://doi.org/10.3390/s21237990

Rego, J. L., & Li, C. (2010). Nonlinear terms in storm surge predictions: Effect of tide and shelf geometry with case study from hurricane Rita. *Journal of Geophysical Research: Oceans*, *115*(C6). https://doi.org/10.1029/2009jc005285

Sarkar, D., Osborne, M., & Adcock, T. (2016). A machine learning approach to the prediction of tidal currents. In *ISOPE International Ocean and polar engineering conference*. ISOPE-I.

Sarkar, D., Osborne, M. A., & Adcock, T. A. (2018). Prediction of tidal currents using Bayesian machine learning. *Ocean Engineering*, *158*, 221–231. https://doi.org/10.1016/j.oceaneng.2018.03.007

See, L., & Openshaw, S. (1999). Applying soft computing approaches to river level forecasting. *Hydrological Sciences Journal*, *44*(5), 763–778. https://doi.org/10.1080/02626669909492272

See, L., & Openshaw, S. (2000). A hybrid multi-model approach to river level forecasting. *Hydrological Sciences Journal*, *45*(4), 523–536. https://doi.org/10.1080/02626660009492354

Stammer, D., Ray, R., Andersen, O. B., Arbic, B., Bosch, W., Carrère, L., et al. (2014). Accuracy assessment of global barotropic ocean tide models. *Reviews of Geophysics*, *52*(3), 243–282. https://doi.org/10.1002/2014rg000450

Tan, E., Algar, S., Corrêa, D., Small, M., Stemler, T., & Walker, D. (2023). Selecting embedding delays: An overview of embedding techniques and a new method using persistent homology. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *33*(3). https://doi.org/10.1063/5.0137223

Tebaldi, C., Strauss, B. H., & Zervas, C. E. (2012). Modelling sea level rise impacts on storm surges along US coasts. *Environmental Research Letters*, *7*(1), 014032. https://doi.org/10.1088/1748-9326/7/1/014032

Wedler, M., Stender, M., Klein, M., Ehlers, S., & Hoffmann, N. (2022). Surface similarity parameter: A new machine learning loss metric for oscillatory spatio-temporal data. *Neural Networks*, *156*, 123–134. https://doi.org/10.1016/j.neunet.2022.09.023

Wells, J. T. (1995). Tide-dominated estuaries and tidal rivers. In *Developments in sedimentology* (Vol. 53, pp. 179–205). Elsevier. https://doi.org/10.1016/s0070-4571(05)80026-3

Williams, J., Irazoqui Apecechea, M., Saulter, A., & Horsburgh, K. J. (2018). Storm surge forecasting: Quantifying errors arising from the double-counting of radiational tides. *Ocean Science Discussions*, *2018*, 1–21.

Williams, J., Saulter, A., O'Neill, C., Brown, J., & Horsburgh, K. (2018). A reassessment of the UK operational surge forecasting procedure.

Woodworth, P. L., & Vassie, J. M. (2022). Reanalyses of Maskelyne's tidal data at St. Helena in 1761. *Earth System Science Data*, *14*(9), 4387–4396. https://doi.org/10.5194/essd-14-4387-2022

Wray, J., & Green, G. G. (1994). Calculation of the volterra kernels of non-linear dynamic systems using an artificial neural network. *Biological Cybernetics*, *71*(3), 187–195. https://doi.org/10.1007/s004220050081

Xie, W., Xu, G., Zhang, H., & Dong, C. (2023). Developing a deep learning-based storm surge forecasting model. *Ocean Modelling*, *182*, 102179. https://doi.org/10.1016/j.ocemod.2023.102179