JGR Oceans

RESEARCH ARTICLE

10.1029/2025JC022758

Response-Based Prediction of Tidal Currents

Thomas Monahan[1], **Tianning Tang**[1,2], **Stephen Roberts**[1], **and Thomas A. A. Adcock**[1]

[1]Department of Engineering Science, University of Oxford, Oxford, UK, [2]Department of Mechanical and Aerospace Engineering, University of Manchester, Manchester, UK

Key Points:

- Explicitly coupling the orthogonal tidal current velocities in the response method enables more accurate predictions
- Response method improves estimates of harmonic constituents from short and noisy data-series
- RTide's automated, data-driven analysis procedure reduces expertise required to conduct response analyses

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). Response-based prediction of tidal currents. *Journal of Geophysical Research: Oceans*, *130*, e2025JC022758. https://doi.org/10.1029/2025JC022758

Received 16 APR 2025
Accepted 8 DEC 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
Project administration: Thomas Monahan, Thomas A. A. Adcock
Resources: Thomas Monahan, Thomas A. A. Adcock
Software: Thomas Monahan
Supervision: Tianning Tang, Stephen Roberts, Thomas A. A. Adcock
Validation: Thomas Monahan
Visualization: Thomas Monahan
Writing – original draft: Thomas Monahan

Abstract This study evaluates the response method for predicting tidal currents. We introduce a coupled response model which explicitly accounts for interactions between velocity components. By leveraging non-parametric and data-driven weight estimation, the approach demonstrates superior predictive accuracy compared to classical harmonic analysis (HA), particularly for fast-moving and non-linear tidal currents. Using ADCP data from the world's largest deployment of tidal stream turbines, the coupled model achieves superior accuracy with fewer than 30 days of input measurements compared to HA using over 180 days of data. Accuracy improvements extend to both current predictions and the derived harmonic constituents, obtained through a specialized procedure. The response approach shows greater robustness when applied to extremely sparse data. This is reflected by the pseudo-admittances, which also show the non-parametric approach advanced can effectively capture unsmooth deviations in the admittance. Analysis of 40 active NOAA current stations highlight when the response approach should and should not be used, yielding average reductions in absolute error of 9.6%. The framework offers new opportunities for studying non-tidal forcing and sediment transport and has significant implications for tidal energy site development. The proposed method is implemented in the open-source RTide Python package, providing a practical and accessible tool that reduces the level of expertise required to apply the response method to higher-order nonlinear processes.

Plain Language Summary This study looks at a new way to predict tidal currents. We develop a method that better captures how different parts of the current interact with each other. Instead of relying on predefining these interactions, it uses patterns learned directly from data. This makes it more accurate than the traditional approach, especially in places where tides move quickly or behave unpredictably. We tested the method using data from the world's largest tidal stream energy site and found that it could deliver better predictions using just 30 days of data—while the older method needed six times as much. Across 40 other NOAA sites, the new method also reduced prediction errors by nearly 10%. This improved approach could help us better understand how tides interact with other forces like wind and sediment, and it could make a real difference in planning and operating tidal energy projects. It is available for anyone to use through the free, open-source RTide Python package—and importantly, it works without needing expert knowledge to set it up.

1. Introduction

Tidal currents present unique challenges for tidal analysis and prediction (Parker, 2007). These tasks are important for a range of engineering and operational applications including tidal energy resource assessment, offshore wind scour prediction, and maritime safety to name a few (Blunden & Bahaj, 2007; Monahan et al., 2023; Schendel et al., 2018). Unlike tidal elevations, tidal currents exhibit significant vertical and horizontal variability (Prandle, 1982). In shallow-waters these flows can be strongly nonlinear, far more so than water levels, and thus exhibit large proportions of energy at higher tidal frequencies. Such complexities are exacerbated by inertial effects which contribute to further nonlinearity and asymmetry (Prandle, 1997).

Classically, analysis and prediction of tidal currents are conducted with harmonic analysis (HA) (Parker, 2007). As with tidal levels, the flow velocity can be decomposed into a set of waves, termed harmonic constituents, which are *assumed* to exist at frequencies identical to the periodic motions of the Earth, Moon, and Sun. For tidal heights, the tidal elevation $\hat{\zeta}$ is given by the superposition of these waves such that

$$\hat{\zeta} = \sum^{N} A_k \sin \omega_k t + B_k \cos \omega_k t, \tag{1}$$

Writing – review & editing:
Thomas Monahan, Tianning Tang, Stephen Roberts, Thomas A. A. Adcock

where A and B are the quadrature amplitudes which correspond to the k^{th} constituent. The amplitude and phase of the harmonic constituents are given by $C_k = \sqrt{A_k^2 + B_k^2}$ and $\theta_k = \arctan A_k/B_k$. Nonlinearity is introduced by including frequencies produced by the sums and differences of these linear constituents. Modern HA is performed in the time-domain, typically utilizing a least-squares type fitting, though other estimators have been introduced to improve the robustness of this approach (Codiga, 2011; Leffler & Jay, 2009; Monahan et al., 2025c). Tidal currents are predicted in this framework by defining constituent ellipses given by pairs of constituents corresponding to the orthogonal velocities $\vec{u}$ and $\vec{v}$. In this work, $\vec{u}$ and $\vec{v}$ refer to the velocities in the East–West and North–South directions respectively. Whilst tidal currents are described by these ellipses, it is important to note that the estimated orthogonal constituents are treated independently. The selection of which constituents to include in an analysis is critical to the success of a HA (Monahan et al., 2025c). Typically, constituent selection is done using the Rayleigh Criterion which defines the "separability" of two constituents with frequencies ω_1 and ω_2 based on the time-series duration T such that

$$\left(\frac{T}{\omega_1 - \omega_2}\right)/R_{\min} \geq 1, \tag{2}$$

with $R_{\min}$ governing the minimum allowable frequency separation for inclusion and is often chosen to be 1. In reality, the extent to which closely spaced frequencies can be separated is limited by the time-series length and the noise content of the signal (Munk & Hasselmann, 1964). While more advanced constituent selection procedures have been developed (Lobo et al., 2024), HA in this manuscript is carried out using the Rayleigh criterion as that is almost exclusively what is utilized by industry. Our proposed approach avoids the pre-definition of frequencies altogether and instead utilizes the data itself to inform the model structure and derived constituents as will be discussed.

In addition to the complex physical characteristics of currents, data themselves present further challenges. Current measurement devices such as Acoustic Doppler Current Profilers (ADCP) are considerably more expensive and difficult to maintain than their water-level counterparts. These elevated operational costs often lead to tidal current records rarely exceeding 3 months (Stiven et al., 2011). This limits the number of harmonic tidal constituents which can be estimated and thus the accuracy of the predictions (Munk & Hasselmann, 1964). These limitations create problems for the development of tidal energy sites which rely on accurate assessments of the stationary tidal current velocity to inform optimal site design (Monahan et al., 2023).

With these operational challenges in mind, it is worth remarking from the outset what this paper will address. This paper describes a method for predicting tidal currents for operational and engineering contexts. While the model can be analyzed for scientific research, prediction is our priority in conjunction with ease of use. The design of the methodology is thus inline with these objectives. A key feature of an easy to use method is also understanding when you should not apply it, which we discuss at length in Sections 4 and 5. There have been extensive contributions to the tidal currents literature which devise methods for analyzing these signals and separating out nonstationary components (Flinchem & Jay, 2000; Lobo et al., 2024; Matte et al., 2013; Pan et al., 2018). These works are important, as the accuracy of the derived stationary signal (both for harmonic and our proposed method) is hampered by noise introduced by transient non-tidal processes such as wave-action, turbulence, changing stratification, or in estuarine locations, fluvial contributions (Jay & Flinchem, 1997; Jay & Musiak, 1996). As such, we advocate our method be used in conjunction with these approaches rather than to replace them. These tools, especially wavelet type analyses, can help to perform an initial site assessment to evaluate the degree to which the observed signal can be regarded as stationary. Consideration of how the proposed approach can be expanded to include non-tidal forcing, and the intrinsic advantages for doing so over a conventional harmonic method, is provided in the discussion and illustrated in Appendix C.

Due to the complexities introduced by nonstationary processes, the extent to which the stationary component of tidal currents can be predicted in the long-term remains an open debate (Parker, 2007; Prandle, 1982, 1997). Here, we argue the limits of this predictability are at least, in part, imposed by the harmonic method itself. To overcome this, we develop a modified response method, similar to the method devised in Monahan et al. (2025a). Unlike the classical response method (Munk & Cartwright, 1966), this approach makes the estimation of nonlinearities simple and automated. Rather than treating the orthogonal current components as independent as in harmonic

analysis, we develop a coupled approach. This allows the model to explicitly learn the coupled effects between $\vec{u}$ and $\vec{v}$, which is not possible in a harmonic analysis. We rigorously show the proposed approach yields superior predictions across virtually any forecast horizon across several performance metrics, but often requires less data to achieve this. These performance improvements are quantifiable within the classic harmonic framework by transforming the response model into a harmonic equivalent model (See Section 2.7).

The remainder of the manuscript is as follows. First, we introduce the response formalism and our coupled model of tidal currents. Next, response and harmonic methods are evaluated on simulated tidal currents in the Pentland Firth, UK. We then evaluate the efficacy of each approach for tidal energy site-assessment using real-data from the largest tidal stream energy site in the world. Finally, both methods are compared across 40 active NOAA current stations.

2. Methods

2.1. Response Method

The response method exploits the fact that the oceanic response to forcing is weakly nonlinear and time-invariant (Munk & Cartwright, 1966). In the case where limited non-tidal forcing exists, the response to the gravitational potential $V(t)$ produced by the Moon and Sun can also be approximated as weakly nonlinear and time-invariant. We here consider applications where this condition holds; however, the approach is readily generalizable to arbitrary forcing as described in Section 5. The response can be fully characterized by a finite set of past, present, and sometimes future values of $V(t)$. To account for this, the tidal potential is expanded in spherical harmonics such that

$$V(\theta,\lambda;t) = g\sum_{n=0}^{n}\sum_{m=0}^{n}\left[a_n^m(t)G_n^m(\theta,\lambda) + b_n^m(t)H_n^m(\theta,\lambda)\right], \tag{3}$$

where g is the gravitational constant, θ is the Geographical Colatitude, λ the East longitude, G_n^m and iH_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 complex spherical harmonics $G_n^m + iH_n^m$ converge rapidly and are defined as

$$G_n^m + iH_n^m = (-1)^m\left[\frac{2n+1}{4\pi}\right]^{\frac{1}{2}}\left[\frac{(n-m)!}{(n+m)!}\right]^{\frac{1}{2}}P_n^m(\cos\theta)e^{im\lambda}, \tag{4}$$

so that the tidal species of interest is governed by the selection of m. Here, P_n^m is the associated Legendre function. Gravitational inputs are restricted to $n = 2$ and $n = 3$ only. While the radiational potential, introduced in (Munk & Cartwright, 1966), was found useful to deal with seasonal effects, we find it degrades performance when performing long-term predictions and is thus not considered further.

The linear response prediction $\hat{\zeta}$ is given by convolving a set of learned weights $w_n^m = x_n^m + iy_n^m$ with lags $[\tau_1,\tau_2,\ldots,\tau_s]$ of the expanded complex gravitational potential $c_n^m = a_n^m + ib_n^m$ such that

$$\hat{\zeta}(t) = \sum_{m,n}\sum_{s}\left[x_n^m(s)a_n^m(t-\tau_s) + y_n^m(s)b_n^m(t-\tau_s)\right]. \tag{5}$$

The learned model $\hat{\zeta}$, can be viewed as a solution to the Laplace Tidal equations after integrating out the real ocean topography.

Similar to the harmonic method, nonlinearity is introduced by taking sums and products of the linear response. The i^{th} order response, with degree and order m,n, takes the form of

$$i^{th}\text{order response} = \sum_{i}\cdots\sum_{x}\sum_{s}\cdots\sum_{s'}w(i,\ldots,x,s,\ldots,s')\,(c(t-\tau_s))(\ldots)(\ldots)(c(t-\tau_{s'})). \tag{6}$$

It can be seen that the number of possible terms grows exponentially with the order of the nonlinearity. As a consequence, only a subset of possible terms are used. Previous works have instead adopted a sequential procedure which begins by first obtaining a linear prediction $\hat{\zeta}$ as in Equation 5, and then forming products of the linear terms $\hat{\zeta}_m^n$ and $\hat{\zeta}_{m'}^{n'}$ of species 1 and 2 to serve as input into a secondary response analysis. This procedure has the flaw of assuming that all non-linearity is locally generated which can lead to the analysis underestimating the variance for higher species. Another challenge is that these terms must be pre-selected which is only achieved through a combination of experience and preliminary spectral analysis (D. Cartwright, 1982). As noted in Monahan et al. (2025b), this limitation has inhibited the automated application of the response method. Both deficiencies are addressed by the non-parametric approach which is introduced in Section 2.4.

Despite the difficulty of estimating nonlinear terms in the classic approach, the response formalism possesses several intrinsic advantages. First, numerous studies have found the response approach to account for slightly more variance than conventional HA using fewer parameters (Cartwright, 1968, 1982; Cartwright & Ray, 1990; Munk & Cartwright, 1966). Second, nodal corrections are implicit (e.g., embedded in the forcing itself), which greatly simplifies long-term tidal analysis and prediction tasks (Munk & Cartwright, 1966). Exactly as the response method relates the tidal potential to the observed tidal variability, the learned response function explicitly models departures from equilibrium nodal theory which is necessary in many regions globally (Woodworth, 2012). A potential difficulty encountered here is that when applying this approach to short-reference series, the "nodal" contribution will not necessarily be constrained. For routine harmonic analysis, the nodal contribution can be toggled off in these contexts. While no such feature exists in a response approach, this feature can be exploited when recovering the "harmonic equivalent model" described in Section 2.7. This paper does not consider reference series long-enough to evaluate nodal effects. Investigation into how nodal effects, especially non-equilibrium effects, warrants future work. Lastly, the response analysis explicitly separates out the oceanic response to different types of forcing, enabling mechanistic studies and the ability to account for other influences such as meteorological and fluvial effects. In order to achieve this, these forces need to be included as additional inputs. In principle, this can be done in the same way as radiational forcing in the original method; however, care needs to be taken to ensure the forcing is adequately represented. With regard to gravitational forcing, the computation of the tidal potential is implemented exactly as in (Munk & Cartwright, 1966). Relevant astronomical arguments are computed directly via the Skyfield Python library with a positional accuracy within 0.0005 arcseconds (Rhodes, 2019). By avoiding a time-harmonic expansion, our program does not need to be redefined for different epochs. A trade-off is that this limits the ability to translate easily between a response and harmonic method. A procedure for doing so it described in Section 2.7.

2.2. Response Tidal Currents

As in the harmonic method, the application of the response method to tidal currents can simply be performed on each orthogonal velocity time-series independently. From these analyses, response weight ellipses can be obtained, and if desired, response admittance "ellipsoids" could be produced.

Here, we propose a "coupled" response model of tidal currents. As will be shown, this approach reduces to the classical separable form under certain conditions, but critically allows for a coupling between the $\vec{u}$ and $\vec{v}$ components. This intuition comes directly from the shallow-water equations in which the orthogonal velocities are directly coupled via both the Coriolis force and advective terms which gives rise to the Bernoulli effect. This coupling can create complex relationships between the orthogonal components at both the same and different frequencies. Regardless, the coupled response model for two orthogonal components of tidal currents can be written as

$$\hat{\beta}(t) = \begin{bmatrix} \hat{u}(t) \\ \hat{v}(t) \end{bmatrix} = \sum_{m,n} \sum_{s} W_n^m(s)\, c_n^m(t - \tau_s), \tag{7}$$

where the weights W_n^m correspond to the matrix

$$W_n^m(s) = \begin{bmatrix} x^m_{n,\vec{x}}(s) + y^m_{n,\vec{x}}(s) & x^m_{n,\vec{x}\to\vec{y}}(s) + y^m_{n,\vec{x}\to\vec{y}}(s) \\ x^m_{n,\vec{y}\to\vec{x}}(s) + y^m_{n,\vec{y}\to\vec{x}}(s) & x^m_{n,\vec{y}}(s) + y^m_{n,\vec{y}}(s) \end{bmatrix}, \tag{8}$$

here the diagonal terms represent the individual sea-level response to the input forcing for the $\vec{u}$ and $\vec{v}$ directions respectively. The off-diagonal terms capture the interactions between them. In the limit where $\vec{u}$ and $\vec{v}$ are independent (as assumed by HA), the matrix $W_n^m(s)$ is diagonal. Section 3.3 provides further evidence for this innovation and considers the improvements it garners in predictive accuracy. Different coupling mechanisms will couple the $\vec{u}$ and $\vec{v}$ components in different ways. As such, a difficulty of this approach is that it makes the selection of which terms to include in an analysis even more complex than in the classic method.

2.3. Lag Selection

Conventionally response analysis has been carried out using uniformly spaced lags where $\tau = 2$ days. This choice has natural advantages when performing the classical spectral analysis approach as the complete linear response can be described by the Fourier transform (admittance function) of the response weights. Indeed, in this framework the choice of lag spacing $\Delta\tau$ has implications regarding the "wiggliness" of the associated admittance. This observation gave rise to Munk and Cartwright's credo of smoothness which found no observations of "wiggles" less than $1/6$ cycles per day (W. H. Munk & Cartwright, 1966).

As in Monahan et al. (2025b), we utilize non-uniform spaced lags based on the quarter period $2\pi/\omega_i$ of the tidal frequencies which can be estimated from an 8-day reference series (M2, K1, M4, M6, M3, M8, 3MK7, 2SK5, 2MK5). In addition to yielding superior performance, this procedure yields a maximum lag of 6.25 days which nicely conforms to the original credo of smoothness (Munk & Cartwright, 1966).

2.4. Weight Estimation

It remains to be said how the response weights w are estimated. In the original response method, weights are estimated via spectral analysis independently for each orthogonal velocity component. Let q be an arbitrary velocity direction ($\vec{u}$ or $\vec{v}$). Here, the frequency spectra H and G are computed for the observed velocity q^j and a given input function γ^j respectively such that

$$H^j(f) = \int_{-\infty}^{\infty} q^j(t) e^{2\pi i f t} \mathrm{d}t, \tag{9}$$

and

$$G^j(f) = \int_{-\infty}^{\infty} \gamma^j(t) e^{2\pi i f t} \mathrm{d}t. \tag{10}$$

These two terms are related to the complex weights w_n through the response relation

$$\hat{q}^j(t) = \int_0^{\infty} \gamma^j(t - \tau_s)\, w_n^j(s) \mathrm{d}\tau. \tag{11}$$

In the frequency domain, this is equivalent to

$$H^j(f) = Z^j(f) G^j(f) \tag{12}$$

where $Z_j(f)$ is the Fourier transform of the response weights w_n^j given by

$$Z^j(f) = \int_0^{\infty} w_n^j(\tau) \exp(2\pi i f \tau) \mathrm{d}\tau. \tag{13}$$

This is the tidal *admittance* function (Zetler & Munk, 1975). As noted by Munk and Cartwright, $Z^j(f)$, actually provides more insight than the response weights w themselves. If discrete estimates of $H^j(f)$ and $G^j(f)$ can be computed, the admittance function is readily obtained through cross-spectral analysis as in (Cartwright, 1968; Munk & Cartwright, 1966). The major shortcoming of this approach is that time-series must be uniformly sampled without gaps—something which is rarely encountered for real ADCP deployments. Additionally, in reality time-series are not infinite leading to frequency resolution limitations. This leads to a smoothing of the response function in the low-data limit which can limit information about long-term variations in tides.

One can also perform a least-squares type procedure as in modern HA. This can mitigate some of the frequency resolution issues by working in the time-domain, but will always encounter similar challenges. Here, the response prediction $\hat{q}$ can be cast as a linear model with $\hat{q} = wx_i + \epsilon$. The design matrix x is composed of the different spherical harmonic input functions at the associated orders, degrees and time lags. While this procedure can be applied to non-uniformly spaced reference series, complications arise when handling areas with considerable nonlinearity. This is a consequence of having to define these interactions a priori which becomes exponentially more difficult as the degree of nonlinearity increases. Unlike the Rayleigh criterion in HA, no such approach exists for the automated determination of nonlinear interactions in a response method.

Motivated by the limitations of the above two approaches, this manuscript advocates an alternate non-parametric data-driven approach originally devised in Monahan et al. (2025a) and implemented in the RTide Python package. Our approach, exploits an equivalence between fully connected neural networks and the Volterra series—the mathematical basis of the response method (Marmarelis & Zhao, 1997; Wray & Green, 1994). Specifically, we utilize a three-layer feed forward neural network, with layers of size equal to the number of inputs (e.g., input functions at each time-lag). This choice was found to effectively capture the system memory for a wide range of tidal systems whilst offering sufficient complexity to represent complex and multiscale phenomena.

Rather than predefine interactions between input functions, the network learns the relationships between inputs simultaneously with the weights, directly from the data. To be explicit, each spherical harmonic input function V_i at time-lag τ is treated as a distinct feature and is fed directly into the network. The weight estimation is performed iteratively using back-propagation and gradient descent based on a given loss-function. Updating the weights in this manner allows the model to iteratively learn the response weights, and nature of the interactions between features directly from the data. Once the weights have been learned, the response prediction, defined by Equation 7, can be replaced by

$$\hat{\beta}(t) = \begin{bmatrix} \hat{u}(t) \\ \hat{v}(t) \end{bmatrix} = f(V_0(t-\tau_0), \ldots, V_0(t-\tau_S), \ldots, V_X(t-\tau_0), \ldots, V_X(t-\tau_S)), \tag{14}$$

where $f(\cdot)$ is the neural network approximation of the time-invariant impulse-response function and V_i are the set of all input functions at the associated time lags. The weights associated with f govern the weights and biases of the hidden layers of the network and are thus not interpretable in the way the least-squares approach is. To overcome this, we propose transforming the model back to its harmonic equivalent which is outlined in Section 2.7.

By definition, this approach gives up the ability to specify which interactions the model should look for a priori. Instead, the model can be viewed in a Bayesian context as only learning interactions for which there is sufficient evidence in the data to support. Here, the tidal potential behaves as a prior for our model, yielding improved robustness in low-data scenarios which will be discussed in Section 3.2 and in Figures S4 and S5 of Supporting Information S1. However, if sufficient evidence exists, the model can also account for strong deviations from a smooth admittance (See Appendix B). Methods for ensuring the model does not overfit to observed data are given below. For those conducting tidal research, this decision may be limiting. However, for those interested in operational or engineering decision making and lacking extensive tidal expertise, this framework enables the usage of a response approach for the first time. We discuss how users can assess whether the response approach is appropriate for their data in Sections 4 and 5.

2.5. Loss Function

The training of the neural network based response model is an iterative process which relies on an optimization objective, or *loss function*. In a classical least-squares harmonic analysis the loss function is a simple mean-squared error. In all contexts, the choice of loss-function has implications for the accuracy of the learned model, particularly when dealing with noisy data as it can be seen to change the statistical assumptions about the data. Similar to harmonic analysis, by default, RTide uses the mean squared error (MSE):

$$L_{\mathrm{MSE}}(\beta,\hat{\beta}) = \frac{1}{n}\sum_{i=1}^{n}(\beta-\hat{\beta}_i)^2.$$

This can be seen to place a greater weighting on outliers which can lead to the biasing of the learned model. Similar to what (Codiga, 2011; Leffler & Jay, 2009) have utilized for least-squares estimation, we utilize the Huber loss which transitions from a quadratic penalty to a linear one for values above a given threshold δ.

$$L_{\mathrm{Huber}}(\beta,\hat{\beta}) = \begin{cases} \frac{1}{n}\sum_{i=1}^{n}\frac{1}{2}(\beta-\hat{\beta}_i)^2, & \text{if } |\beta-\hat{\beta}_i| \le \delta \\ \frac{1}{n}\sum_{i=1}^{n}\delta\left(|\beta-\hat{\beta}_i| - \frac{\delta}{2}\right), & \text{if } |\beta-\hat{\beta}_i| > \delta \end{cases}$$

By default, we utilize $\delta = 1$. This was found to improve performance for the operational applications in Section 4 which exhibit significant noise. Since the data are normalized during training (mean removed and scaled to unit variance), we find this choice of δ to be appropriate for a wide range of scenarios. For simulation data, the mean-squared error loss is used, and for real data, we utilize the Huber loss.

2.6. Architecture

In its simplest form, such as a single-layer network with no hidden nodes, a neural network is equivalent to conventional linear regression. However, by increasing the number of layers, nodes, and activation functions, neural networks can represent far more complex and nonlinear relationships between inputs and outputs. In this context, the *architecture* refers to the overall structure of the neural network, including the number of layers, the number of nodes in each layer, and how they are connected. *Activation functions* introduce nonlinearity, allowing the network to capture complex, curved relationships rather than just straight-line ones. They act like switches that decide when each node becomes active, enabling the model to learn richer patterns in the data. Selection of this structure therefore impacts the nature of the physical interactions learned by the network. This manuscript considers three separate architectures. All architectures receive the same inputs, namely the gravitational spherical harmonics input functions V_i at associated time-lags $\tau_1,\ldots,\tau_S$. A visualization of the different model architectures is given in Figure 1.

2.6.1. Separable RTide

As with HA, one can perform a conventional response analysis on the $\vec{u}$ and $\vec{v}$ components separately. Here, we use an identical architecture to that proposed in Monahan et al. (2025a). The given architecture has 3 hidden layers with the number of neurons in each layer equal to the number of inputs. The final output is simply the summation of the outputs from prior layers. Tanh activation functions are employed to yield equivalence to the Volterra series (Marmarelis & Zhao, 1997). By treating the $\vec{u}$ and $\vec{v}$ components separately, this formulation is equivalent to when W is diagonal in Equation 8.

2.6.2. Coupled RTide

A coupled RTide model is readily constructed by simple modification of the output layer. Here, the final summation layer is replaced by a linear layer with two outputs. Through this simple modification, the model simultaneously predicts and regresses the $\vec{u}$ and $\vec{v}$ components using a shared set of weights. This formulation is equivalent to Equation 7.

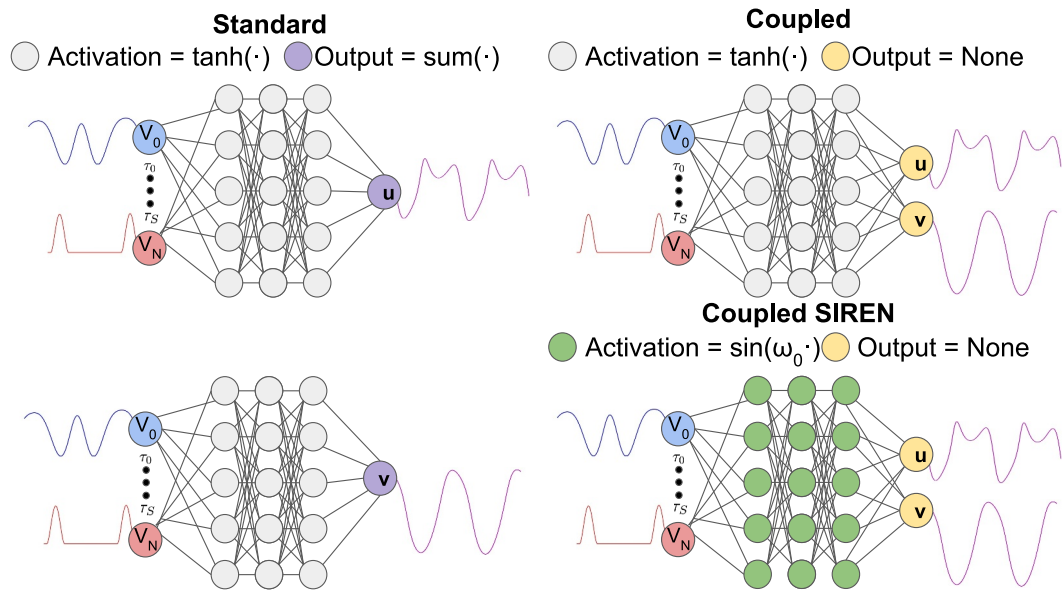

Figure 1. Comparison of RTide model architectures. Activation refers to the activation function used in the hidden layers of the network. Output refers to the activation function used for the output node(s) of the network. If the activation = None, no nonlinearity is applied and thus the final output is simply given by $wx + b$ where x is the input and w and b are the learned weight and bias.

2.6.3. SIREN RTide

Fast-moving tidal currents often exhibit significant fractions of energy at high-frequencies (Parker, 2007). Effectively capturing these multiscale features is difficult. To combat this we replace the Tanh activation functions with sinusoidal ones. This is known as a SIREN architecture and is well-suited to modeling the oscillatory and high-frequency nature of these currents (Sitzmann et al., 2020). While these networks are more expressive, they can also be more difficult to train. The activation is defined as:

$$\text{SineActivation(inputs)} = \sin(\omega_0 \cdot \text{inputs}),$$

where ω_0 is a trainable or fixed frequency scaling parameter that controls the range of frequencies the network can represent. Here, we find $\omega_0 = 30$ to be a suitable choice after considerable sensitivity testing. The weights of the network are initialized using a specialized initializer, given by:

$$\text{Weights} \sim \mathcal{U}\left(-\frac{1}{XS}, \frac{1}{XS}\right),$$

where $\mathcal{U}$ is a uniform distribution with a scale parameter equal to the inverse of the total number of inputs XS (the number of input functions multiplied by the number of lags). This initialization ensures stable gradient propagation during training and is crucial for avoiding issues like vanishing or exploding gradients which are both common problems in other areas of machine learning. SIREN architectures can be utilized in the RTide Python package by setting "architecture = SIREN" when calling Train(). While this architecture can be used for both coupled and uncoupled predictions, here we only consider its application in the coupled model.

2.6.4. Training Particulars

For training, we first randomly partition the observational data into a training and validation set, usually with a ratio of about 10 to 1 training to validation. The validation set is used to evaluate the generalization of the learned model out-of-sample throughout the training process. Unlike most tidal analysis routines which rely on least-squares estimators, our approach is carried out using gradient descent. This is a type of iterative optimization procedure applicable to our nonlinear and multi-layered models which lack analytical weight optimization

solutions. Here, models are trained for up to 500 iterations (epochs). In each epoch, the model updates its internal weights to reduce prediction errors over the training data set through backpropagation and gradient descent. The learning rate, which controls the size of each update step to the model's weights during training, is initially set to 1×10^{-3} and is progressively reduced if the performance on the unseen validation data (loss) does not improve for more than 10 consecutive epochs.

Early stopping (e.g., ending the training loop) occurs if the validation loss plateaus for more than 30 epochs. As has been carried out in various tidal analysis programs (Leffler & Jay, 2009), we utilize L2 regularization to encourage smaller model weights which we find to improve generalization. By default, this value is set to 0.001, but 0.1 was found to yield superior performance for the SIREN architecture. Fine-tuning the regularization can have significant impacts on the associated analysis. If a model appears to be over, or under-training, we generally recommend this is the first parameter that should be tweaked. In this work, all parameters were set to these default settings in the RTide package. The resultant analysis can be replicated using just three lines of code.

2.7. Harmonic Equivalent Response Model

The learned response model can be transformed to its harmonic equivalent through harmonic analysis of the RTide predictions. Here, we make use of the VTide python package which allows for variational Bayesian harmonic analysis to be conducted as proposed in Monahan et al. (2025c). Gravitational input functions are precomputed for a full year at hourly intervals (8,760 samples) and used to generate a synthetic RTide time-series to be analyzed. Given the response model is trained on far less data than 1-year, care must be taken when computing the harmonic equivalent model. Here, we encounter a tricky question: how do we decide which constituents we can and cannot include? The following approach takes inspiration from several ideas in Innocenti et al. (2022).

First, using the estimated amplitude mean μ and standard deviation σ^2 outputted by VTide from a pilot analysis of the RTide prediction, we compute the SNR of each constituent such that

$$SNR = \frac{\mu}{\sigma^2}. \tag{15}$$

Here, we utilize all constituents which are resolvable from a 1-year reference series using a standard Rayleigh criterion of 1.0. As with other tidal analysis programs (Lobo et al., 2024), we automatically include constituents with $SNR \geq 2.0$. Next, we found it useful to threshold parameters with SNR below 0.5 following empirical analysis. While this is not a requirement for usage of the following procedure, it can reduce the computational cost associated with the recursive algorithm. We conducted tests on a shortened reference series from several stations, using the full reference series as a pseudo–ground truth. Constituents with $SNR < 0.5$ were consistently inaccurate or insignificant, whereas those between 0.5 and 2.0 exhibited greater variability in accuracy and significance. This is a consequence of the fact that the variational Bayesian estimator can underestimate uncertainty as discussed in Monahan et al. (2025c); Innocenti et al. (2022). Hence, this intermediate range requires further examination.

We overcome this difficulty by introducing a principled Bayesian criterion for constituent selection based on Bayes factors (Kass & Raftery, 1995). The Bayes factor, BF_{10}, compares the marginal likelihoods of two competing models M_1 and M_0 given data D, and is defined as

$$BF_{10} = \frac{p(D|M_1)}{p(D|M_0)}. \tag{16}$$

A value of $BF_{10} > 1$ indicates evidence in favor of model M_1, while $BF_{10} < 1$ favors model M_0. This provides a rigorous approach for selecting the optimal set of tidal constituents by quantifying the strength of evidence for each candidate model. This can be seen as more principled than the Bayesian Information Criterion which is a cruder approximation of the log of the Bayes factor due to the lack of consideration of the model prior distribution (Kass & Raftery, 1995). The difficulty when working with Bayes factors is the question of how one estimates the marginal likelihood, or the evidence of the data given the model. Typically, the marginal likelihood is intractable

and must be approximated through sampling. However, in the variational Bayesian framework, optimization is carried out by maximizing the evidence lower-bound (ELBO) for some observed data $p(Y)$ defined as

$$ELBO = F(p(\theta|Y), q(\theta|Y)) + \mathrm{KL}(p(\theta|Y), q(\theta|Y)). \tag{17}$$

where $q(\theta|Y)$ is the approximate posterior, and $p(\theta|Y)$ the true posterior. As the name suggests, the ELBO is a lower-bound on the true marginal likelihood and can thus be used directly as an approximate Bayes factor (Roberts & Penny, 2002). Psuedo code for the iterative algorithm is given in Algorithm 1.

Algorithm 1. Bayesian Constituent Selection via ELBO Improvement

```
constituent_list ← Rayleigh criterion constituents from observed data
non_rayleigh_constituents ← Rayleigh criterion constituents from full year\constituent_list
  Run VB model on RTide prediction using constituent_list
best_ELBO ← ELBO_Rayleigh
for each constituent with SNR ≥ 0.5 in non_rayleigh_constituents do
  Run VB model on RTide prediction with constituent_list + constituent
  if ELBO_new / best_ELBO > 1.01 then
    best_ELBO ← ELBO_new
    constituent_list ← constituent_list ∪ {constituent}
  end if
end for
```

An advantage of applying this procedure to 1-year reference series is that it reduces errors arising from the harmonic estimator itself when applied to short reference series (Innocenti et al., 2022; Monahan et al., 2025c). Since the gravitational inputs are precomputed, the added computational cost consists of (a) the forward pass through the network to generate the predictions, and (b) the variational Bayesian harmonic analysis. The former is on the order of milliseconds, and the latter under a minute when considering 60+ harmonics. Comparisons of UTide and RTide harmonic equivalent models are given in Sections 3.2 and 3.3.

An open question raised by reviewers and not addressed by this automated procedure is how one can select which overtide constituents should be considered as candidates to begin with. Here, we only consider tidal lines used in standard analysis programs (e.g., UTide). For users interested in pure tidal research, there is no substitute for pre and post analysis of the data and learned model. This is particularly important for identifying constituents which may not be related to the tidal potential but instead things like friction and warrants careful consideration of the user.

2.8. 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 $R = 1$ unless stated otherwise. While hand-selection of constituents can increase the accuracy of the derived constituents, we opt to compare with Rayleigh as this is what is utilized by most engineering practitioners. 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).

2.9. Simulated Data

The Pentland Firth is characterized by exceptionally fast tidal currents, thus making it a high-interest site for tidal energy development. We use data from the shallow-water flow model validated by (Adcock et al., 2013, 2014) to produce a depth-averaged simulation of the currents at the Pentland Firth. Simulation data is provided for the period from 01-01-2014 00:00 UTC to 01-01-2015 00:00 UTC at a sampling interval of 300 s. The model is forced with constituents K1, K2, M2, MU2, N2, NU2, O1, and S2 at the ocean boundary.

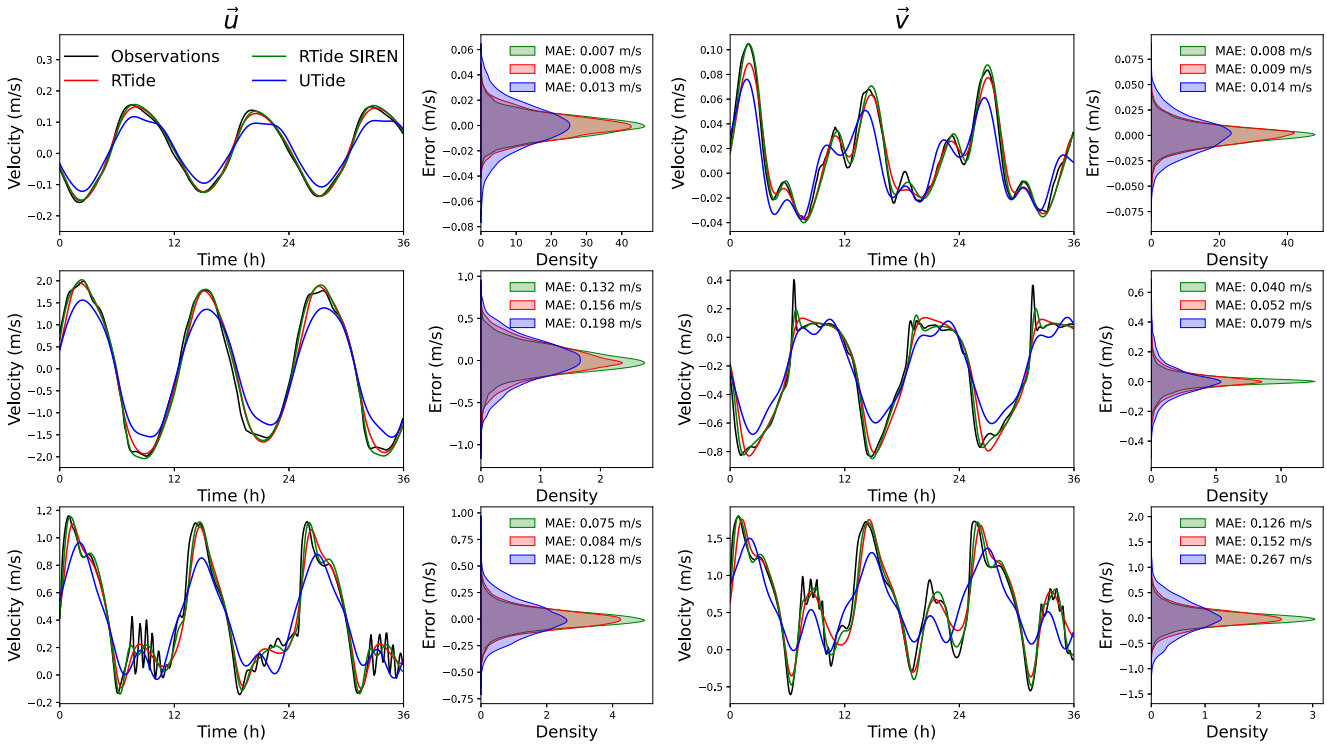

Figure 2. Comparisons of UTide and RTide model's out of sample performance when trained on 180 days of simulated data. Predictions are of the first 1.5 days of the testing data set (04-07-2014 22:40 to 06-07-2014 10:35). Kernel density estimates of the residual error are computed from the subsequent 180 days of test set predictions. Top, middle, and bottom rows correspond to locations A, B, and C respectively with coordinates provided in Table S1 of Supporting Information S1.

3. Validation

3.1. Simulated Comparison

In this section, we test the capabilities of the proposed architectures on simulated data from the Pentland Firth, Scotland described in Section 2.9. Comparisons are made at three candidate sites characterized by increasing amounts of nonlinearity. Note, here we are using the raw numerical model outputs which do not contain instrumental noise errors to assess the model's ability to represent physical processes. We consider the impact of noise in Section 3.2.

Visualizations of the model predictions of the simulated currents are given in Figure 2. Inspection of the kernel density estimate (KDE) of the residuals (computed over the entire 180 days test set) shows both RTide models have a greater density of prediction residuals near zero. This is reflected by the reported mean absolute errors (MAE). Across all three sites, the percent reduction in the Euclidian norm of the MAEs for each model shows a reduction in MAE of 33% and 44% for the coupled RTide and SIREN RTide models respectively.

The reason for this improvement is readily seen in the associated predictions. The tidal velocities exhibit sharp nonlinearities, and significant distortions to the smooth tidal profile. Due to the limited duration of the reference series, HA attempts to fit a *subset* of tidal harmonics to this profile in a least-squares sense. This leads to biasing of the derived constituents, as much of this distortion is brought by the spreading of energy to higher harmonics as will be shown in Section 3.2. In contrast, the RTide models can be seen to capture a majority of this behavior with the improvements over HA increasing with the complexity of the given location. The difference between the standard and SIREN RTide architectures can also be seen at locations B and C. Here, the SIREN RTide predictions not only do a better job of capturing the larger nonlinear distortions, but can also be seen to capture many of the small nonlinear undulations.

3.2. On the Limits of Predictability

Most discussions of tidal current predictability have naturally focused on using harmonic analysis (Godin, 1991; Parker, 2007). While this, of course, makes sense given its ubiquity, any conclusions made cannot be separated from the limitations of the method itself. These reviews have generally painted a pessimistic view regarding the predictability of currents. We here look to distinguish between the limitations imposed by the physical processes, and those of the predictive algorithm.

The reason for drawing such a distinction is readily seen in Figure 3. Here, we consider the frequency spectrum of the $\vec{v}$ and $\vec{u}$ velocities at simulation location C, alongside the associated HA and response predictions. These predictions and the simulated velocities are outside of the models training data and are therefore a true representation of what the model has learned. Several observations can be made. First, significant energy exists at high tidal frequencies, extending beyond 12 cycles per day (not shown for clarity). Based on a Rayleigh threshold of 1, the only constituents included beyond 6 cycles per day are 3MK7 and M8. Due to the spreading of spectral energy in these peaks, HA with just these constituents misses out on considerable variability, not to mention the frequencies beyond 8 cycles per day. It is, of course, possible to include hand-picked additional constituents within the given analysis. However, inspection of the spectral bands beyond 8 cycles per day illustrates why this is so challenging. A naive approach would likely consider adding the overtides M10 and M12, yet these constituents hardly rise beyond the noise-floor in their respective bands. In contrast, the response approach does well to capture a majority of the variations across species and at high frequencies, including the broadening of the spectral peaks. This highlights one of the main advantages of our non-parametric response approach—rather than assume a model a priori, we learn the model directly from the data. For regions with minimal nonlinearity where a linear prediction suffices, this approach can be less efficient. However, as shown, we find that in most cases for tidal currents this approach yields both more accurate predictions from less data and the harmonics of high-frequencies which are missed by HA.

One interesting aspect of a response approach is the associated admittance function. Commonly, tidal inference—that is inference of unresolved minor constituents using known relationships to resolved major constituents—is done by assuming a smooth relationship across narrow frequency bands between major and minor constituents (Le Provost et al., 1991). However, this assumption can break down in areas with strong resonances, or in the case of constituents such as MSN2 which is primarily produced via frictional effects and largely absent in the tidal potential (Le Provost, 1991). In a classic response approach, the selection of the time-lags associated with the analysis limits the "wiggliness" of the admittance function. This is both a strength and a weakness as response approaches can make use of the imposed smoothness across narrow frequency bands, but by definition may fail to capture strong deviations from this assumption. With regard to currents, frictional compound tides such as MSN2 are far larger than equilibrium theory would suggest, generally breaking this assumption. A visualization of the response predicted and true admittance across the semi-diurnal band is shown in Figure B1. To ensure an accurate ground truth, this experiment was conducted by applying a response analysis to a reference series generated using the constituents derived from the synthetic data at location C. Here, the MSN2 constituent has a magnitude of approximately 0.024 m/s (around 2 orders of magnitude smaller than M2). The associated response estimate of the constituent has an RMS error of 0.006 m/s which suggests the model can correctly account for a majority of this variability. This is seen in the admittance in Figure B1, where a majority of the large jump in the admittance is captured. Modeling such deviations in real data will always be dependent on noise; however, this result demonstrates the non-parametric RTide approach is capable of modeling large departures from a smooth admittance given sufficient observational data.

The HA approach can account for more of this variability through the introduction of hand-picked constituents. However, such additions are rarely made by industry practitioners. Furthermore, improving the HA is not as simple as adding more constituents into the analysis. This is a consequence of the relatively high noise-floor seen in $\vec{v}$. Given this is simulation data, this noise stems from physical processes, such as vortex shedding off of islands, and not instrumental noise. Inclusion of the latter would of course make things even more difficult. In this vein, the learned response model can be seen to significantly improve the signal-to-noise ratio around groups of species, and can thus be used as a "denoiser" to obtain superior harmonic predictions. A table of harmonic constituents for both the response and harmonic analysis in Figure 3 is provided in the Supporting Information S1. A true apples-to-apples comparison between HA and response approaches is given in Figure 4.

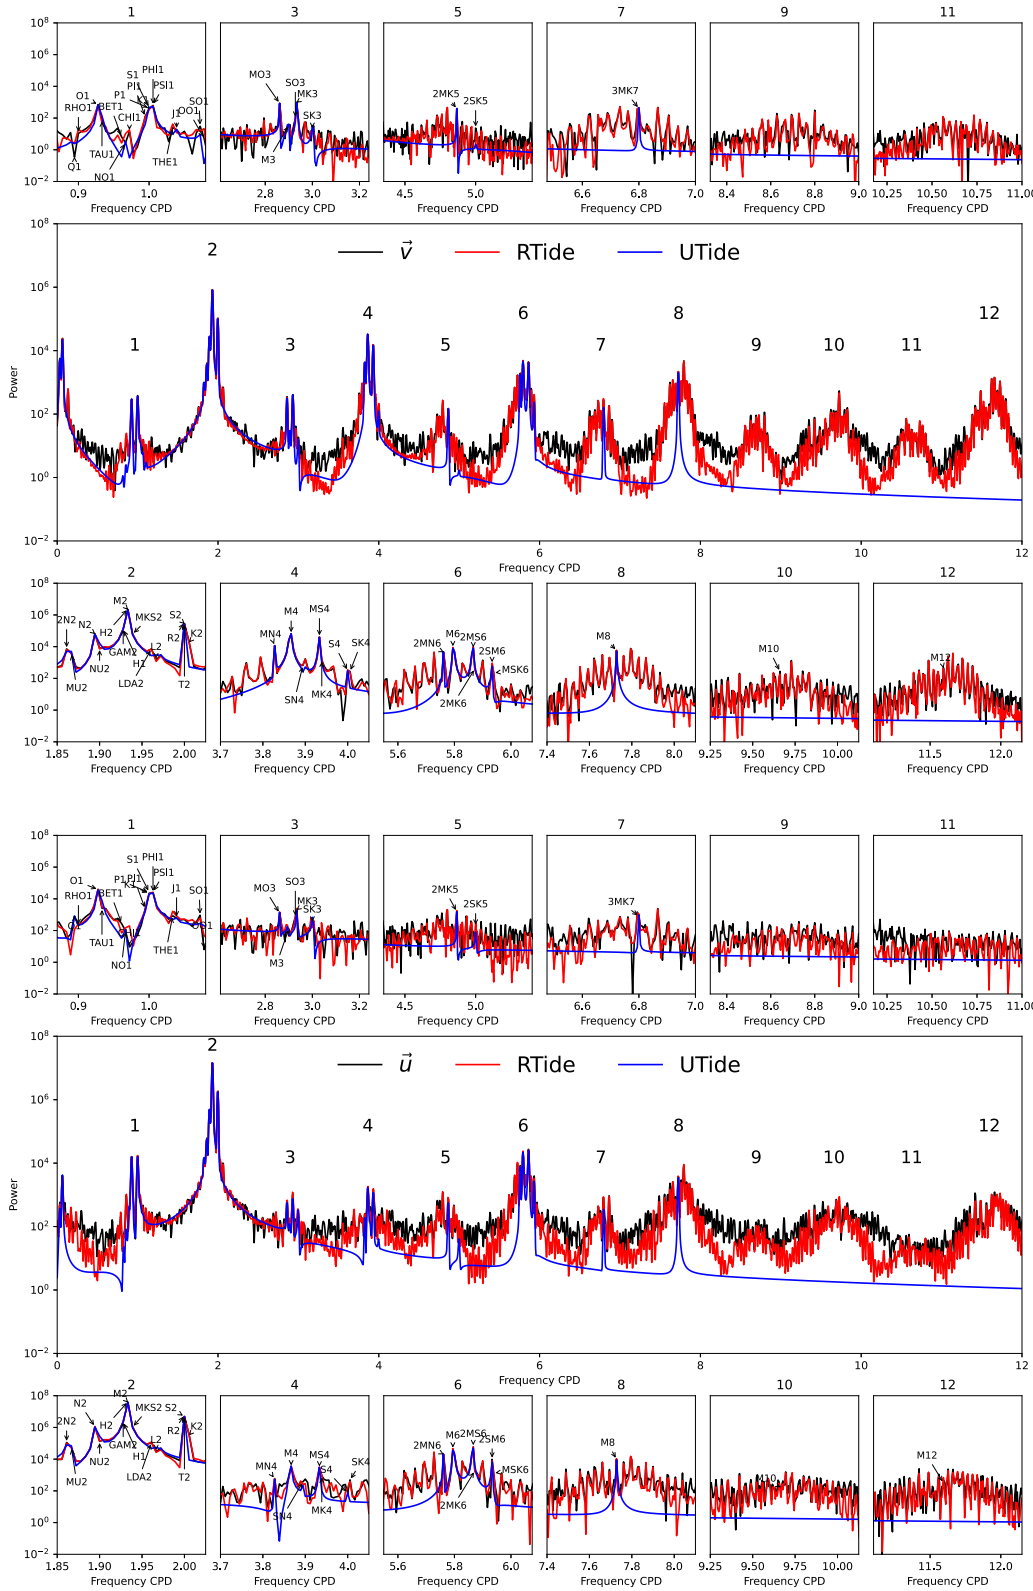

Figure 3. Comparisons of predicted and observed power spectra for the simulated data between RTide and UTide. The top panel shows the *v*-component and the bottom panel shows the *u*-component. Models are trained on 180 days of data, and PSD estimates are obtained for the 180-day test set. Smaller panels are numbered according to the species they correspond to in the central plot. All 67 standard UTide harmonic constituents are labeled.

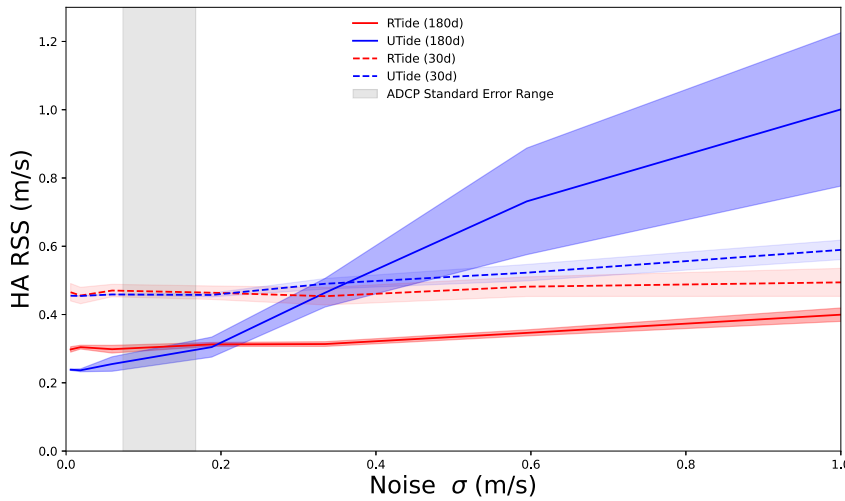

Figure 4. Comparison of RTide and UTide derived harmonic constituent root-sum-of-squares (RSS) error as a function of WGN. To enable a fair comparison, RSS is estimated only for the constituents resolvable using a Rayleigh criterion of 1. Dashed and solid lines correspond to models trained on 30-day and 180-day of data respectively. One standard deviation confidence intervals are estimated from 5 repeat trials. An estimated ADCP standard error range from 3.9% to 8.9% is shown (Klema et al., 2020).

As has long been understood, the ability to resolve individual tidal constituents depends on the noise level in the data (Munk & Hasselmann, 1964). For current measurements, noise levels are typically higher than for water-level records due to the characteristics of the instruments used (Parker, 2007).

Here, we assess how the RTide and UTide models perform in estimating harmonic constituents under varying levels of white Gaussian noise (WGN). The harmonic equivalent response procedure described in Section 2.7 is applied to the Pentland Firth simulation data set. Figure 4 presents the root-sum-of-squares (RSS) error in the estimated constituents as a function of the added WGN amplitude. The RSS is computed as the sum of RMS errors across all constituents resolvable under a Rayleigh criterion of 1.

For both the 30-day and 180-day simulations, the two models exhibit similar behavior. In the low-noise regime (below the ADCP standard error range), pure harmonic analysis yields slightly more accurate results. Within the standard error range (3.9%–8.9% of the velocity standard deviation), the difference between the two methods is negligible. However, as the noise increases beyond this range, their performance diverges. Although WGN does not formally violate the assumptions of least squares, its elevated noise floor limits the resolution of smaller constituents. In contrast, the response-based approach effectively filters this noise, resulting in more robust constituent estimation.

Arguably more problematic are transient non-tidal processes such as wind, storm surge, and changes in stratification which can also temporarily modulate the tidal current. These phenomena are not considered in this manuscript but can introduce problems in analysis. As such, if only gravitational forcing is being considered, it is recommended to conduct a preliminary analysis of the data using either a short-term HA, or a wavelet based analysis (Flinchem & Jay, 2000; Lobo et al., 2024) to assess whether or not that tide can be treated as stationary. Section 4 provides guidelines for identifying when a signal is contaminated by such processes. A scheme for learning the coupled response to these non-stationary processes could also be implemented in RTide as in Monahan et al. (2025a) and shown in Appendix C.

3.3. Tidal Energy Resource Assessment and Prediction

There is considerable economic interest in exploiting the energy of fast-moving tidal flows (Adcock et al., 2013). Due to the significant spatial variability of tidal currents (Prandle & Matthews, 1990), it is necessary to sample multiple locations when conducting tidal energy resource assessments. Here, two problems are encountered. First, as noted above, HA requires an extended reference series to yield adequate accuracy. Second, due to the cost associated with the purchase and maintenance of ADCPs, one must either limit the duration of the measurement

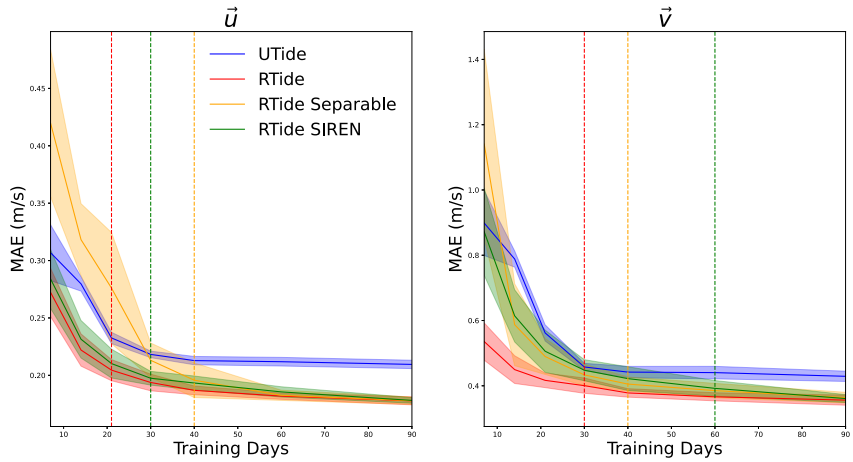

Figure 5. Comparisons of UTide and RTide model's out of sample performance on 2 years of Meygen data when trained on reference series of increasing duration. Solid lines denote mean values over 10 trials on different training data. Errors bars represent one standard deviation. Vertical dashed lines mark where the mean of an RTide model exceeds the performance of UTide trained on 180 days of data.

campaigns at each site, or the number of sites which can be sampled in a given campaign. Naturally, both of these outcomes are unsatisfactory. This section highlights the ability of our modified response-method to reduce the duration of these measurement campaigns whilst maintaining comparable or superior performance.

Current measurements are taken from a turbine-mounted ADCP at the MeyGen Tidal Energy site in the Inner Sound of the Pentland Firth, Scotland (Thomson, 2016). Presently, this is the largest tidal stream turbine deployment in the world. Data is taken from 11-01-2019 00:00 UTC to 28-12-2021 00:00 with measurements sampled at 600-s intervals without smoothing. We compare the performance of Response and HA when trained on data series ranging in length from 7 to 180 days (6 months). Here we compare conventional HA, separable RTide, the new coupled RTide, and the coupled RTide SIREN architecture. To capture the range of model performance when trained at different periods, models are trained over 10 different subsets of the 2 years reference series based on the defined training interval. Mean absolute errors are reported for the out-of-sample test set, which is given by the unused portion of the 2 years of data. The mean and standard deviation of these results are shown in Figure 5.

This analysis reveals several interesting characteristics. First, with regard to the ultra-short training series ($\leq$ 14 days), the coupled RTide model outperforms all other models in terms of MAE. Compared to UTide, this constitutes a 13% and 41% reduction in MAE after only 7-day. The separable RTide model performs the worst of all models for these horizons. Interestingly, while the SIREN RTide performs similarly in terms of MAE to the standard coupled RTide for the $\vec{u}$ velocity (dominant flow direction), it is only marginally better than the UTide for the $\vec{v}$ velocity. The vertical lines indicate where the RTide model exceeds the performance of the UTide model when trained on 180 days. For the standard coupled RTide, this occurs at only 21 and 30 days for the $\vec{u}$ and $\vec{v}$ velocities respectively. The separable and SIREN models are both somewhat less accurate in terms of MAE but still significantly improve on the HA. The reason for this robustness in the low-data regime is in part due to the support provided by the tidal potential. Visualization of the corresponding psuedo-admittances of both standard coupled RTide and UTide are provided in the Figures S4 and S5 of Supporting Information S1. The psuedo-admittances confirm the empirical observations that the coupled RTide approach fails gracefully as the amount of observational data is reduced (even down to 1-day).

As the duration of the data is increased, the UTide MAE plateau higher than the RTide MAEs. The RTide MAEs can also be seen to continue to decrease. This is a consequence of the fact that the non-parametric RTide approaches construct models based on the provided data rather than assuming them beforehand. Due to the parametric nature of HA defined by the Rayleigh criterion, the HA model can only fit the constituents defined at the start of the analysis. As shown in Section 3.2, this procedure misses out on both high-frequency constituents, and those spread across a given species when analyzing shorter reference series. Hence, due to the underlying noise

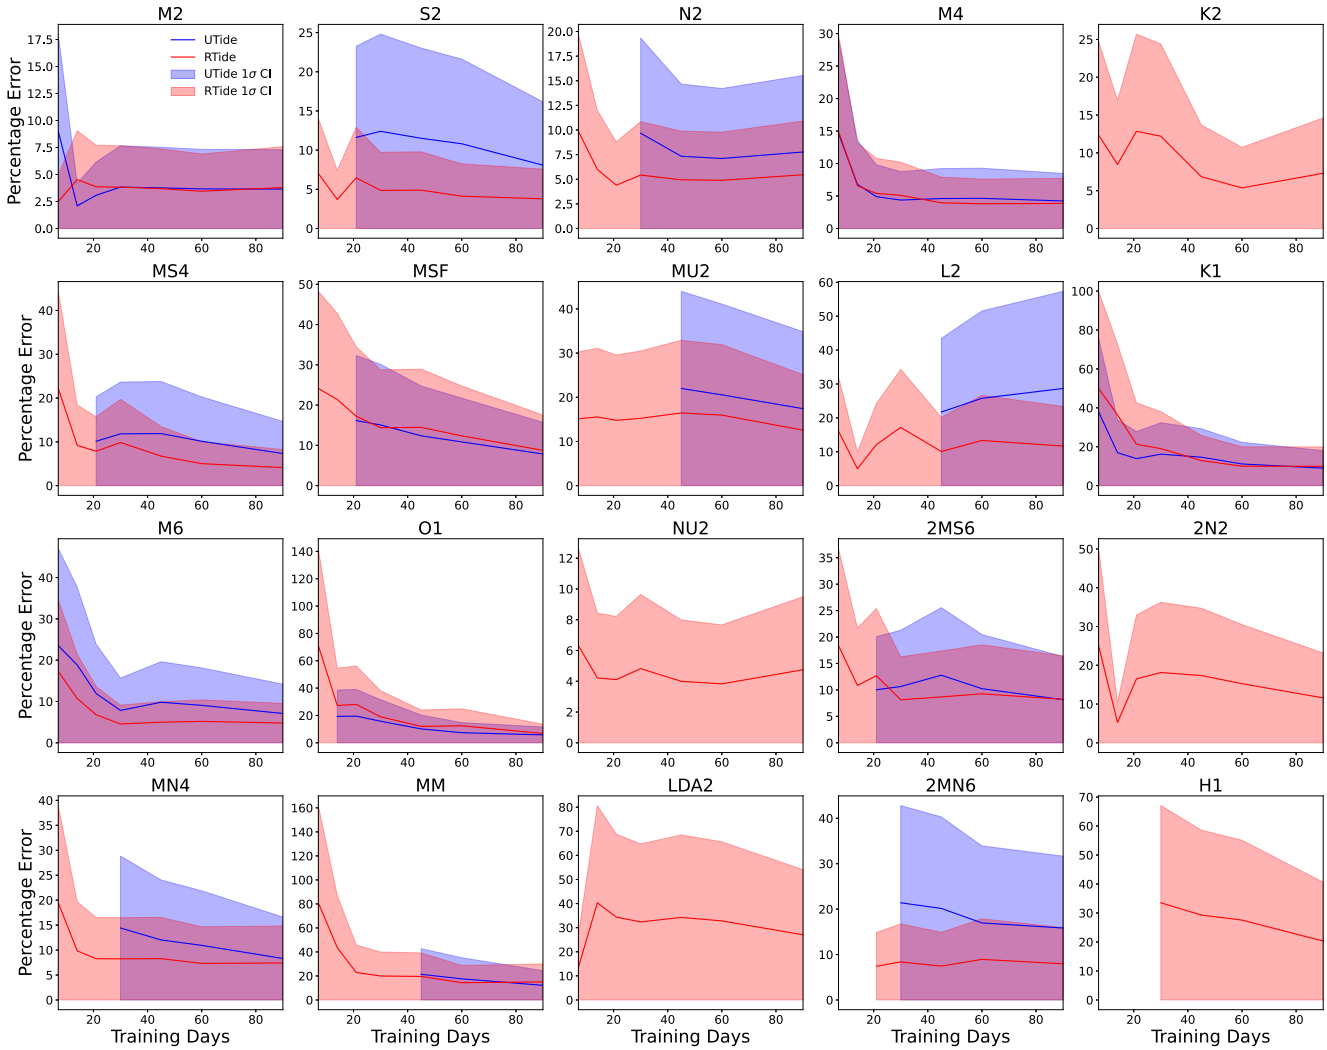

Figure 6. Comparison of RTide and UTide derived harmonic constituent percentage error from Meygen data (Equivalent to Figure 5). Total percentage error is computed as the Euclidian norm of the constituent RMS error for major and minor ellipse parameters. Solid lines denote mean values over 10 trials on different training data. Errors bars represent one standard deviation. UTide constituents are estimated using a Rayleigh criterion of 1. If no estimate exists then insufficient data exists based on this threshold. The "ground-truth" constituents are those derived from the full 2-year Meygen data series.

(both physical and measurement-based), one must wait for more data to introduce more constituents thus leading to the observed behavior.

One can also evaluate the difference between response and harmonic models within the familiar harmonic framework. Harmonically analyzing the predictions of the RTide model allows for a harmonic equivalent RTide model to be produced. Identical comparisons of the UTide and RTide derived harmonic constituents are provided in Figure 6, with the RTide constituent estimates obtained using Algorithm 1. For the RTide estimates, if no value exists at a given training day duration, then either the estimated SNR was too low, or there was insufficient evidence for the constituent given the data in a Bayesian sense. For UTide, the threshold of inclusion is determined using the standard Rayleigh criterion. Here, the constituents derived from the full 2-year of Meygen data are considered as the ground truth. Total percentage errors for the respective constituents are given by the Euclidean norm of the RMS errors of the major and minor ellipse parameters (See Appendix A). Only the 20 largest constituents are shown for clarity. Analysis of individual constituents shows a similar trend to Figure 5 with many of the RTide estimated constituents achieving equal or lower percentage errors using less data than UTide. Some constituents, such as K2 and NU2, which require time-series longer than 210 days using a Rayleigh threshold of 1.0, can be seen to achieve percentage RMS errors under 10% using less than 60 days of data.

However, some problems exist with smaller constituents such as LDA2 and H1. For these small constituents, the RTide models appear to recover the magnitudes and phase of the major axis well (errors less than 1 cm/s and 10° in phase), but struggle with the phase of the smaller minor axis whose magnitude is less than 1 cm/s with errors exceeding 10s of degrees. It can be seen that for H1 and 2MN6, the Bayes factor approach correctly omits the constituent from analyses with fewer that 30 and 21 days of data respectively. The same cannot be said for other problematic constituents such as O1 and MM for reference series shorter than 7-day. A possible, but potentially conservative, mitigation is to perform a posterior predictive check on the observational data with and without the given constituent. This will reflect whether or not there is sufficient evidence in the shortened data series to support the given constituent or not. Future work is needed to understand the frequency resolution characteristics of the response method and to further robustify the constituent selection procedure. Regardless, this case study shows the intrinsic interpretability of HA need not be sacrificed, and can be enhanced, when applying a response based analysis.

4. Operational Forecasting

While the proposed approach yields the greatest improvements for strongly nonlinear tidal currents, we here look to test the model's performance across a wide range of locations. Implicit within this is also understanding where and how the model fails. Response and harmonic models are evaluated on 40 active stations from the NOAA Currents active stations list in the USA. These stations range in total data durations from 13 years to 20 days. Due to the presence of major noise artifacts and sensor movements, we have manually cleaned the data. All 40 NOAA reference series are prepared based on the following criteria. Data is downloaded directly through the NOAA Tides and Currents API with a cut-off date of 09-01-2024. We restrict that all data must be taken from the same location (some sensors can be seen to drift over the years). To avoid this we only train and evaluate models using data from the most recent deployment. We define the end of a deployment as being the first instance in which there is a data gap of greater than 3 days. No further cleaning was done as we wish to test both model's capabilities in handling the noise inherent with real measurements. The available data is split in half, with the first half of observations being used for training, and the second half used for evaluation. The results provided are on the unseen test set.

RTide and UTide models are both trained using the default set of parameters (UTide parameters are described in Section 2.8). While the performance of both methods could certainly be optimized through expert intervention, the present case study aims to provide a true assessment of "off-the-shelf" performance, which was not tractable using traditional response methods. Note that the RTide model employed is the coupled RTide architecture. We found the SIREN architecture to do marginally worse, especially for sites characterized by considerable noise. The change in MAE relative to UTide at each of the 40 stations is shown in Figure 7. The combined MAE at a single station is given by the Euclidean norms of the MAE of the $\hat{u}$ and $\hat{v}$ predictions where

$$MAE_E = \sqrt{MAE(\hat{u})^2 + MAE(\hat{v})^2}.$$

The percentage change in MAE is then computed relative to the combined UTide MAE. It can be seen that RTide yields a reduction in MAE at 31 out of 41 stations. This corresponds to an average reduction in MAE of 9.6%, and a median reduction of 6.5% across all 40 reference stations.

These results agree with those of previous authors which have shown the response method to yield improvements in predictive accuracy over HA for tidal heights (Cartwright, 1968; Munk & Cartwright, 1966; Zetler et al., 1979). In comparison, the average and median reductions of 9.6% and 6.5% across the 40 NOAA current stations are somewhat higher than those found for tidal heights. As shown in Section 3.2, we surmise that these improvements are a consequence of the currents themselves exhibiting greater energy at both higher frequencies and being more spread out within the spectral peaks. In a classical HA, accounting for these components requires a lengthier reference series.

It is worth considering where the response approach is performing worse than the conventional HA. The inset plot within Figure 7 compares the same stations as a function of the ratio of Tidal and Long period energy. To compute this, we use the Lomb-Scargle periodogram to estimate the power spectral density of each velocity component ($\vec{u}$ and $\vec{v}$) separately, as a consequence of the gappy and irregular sampling of some sites. We then integrate the

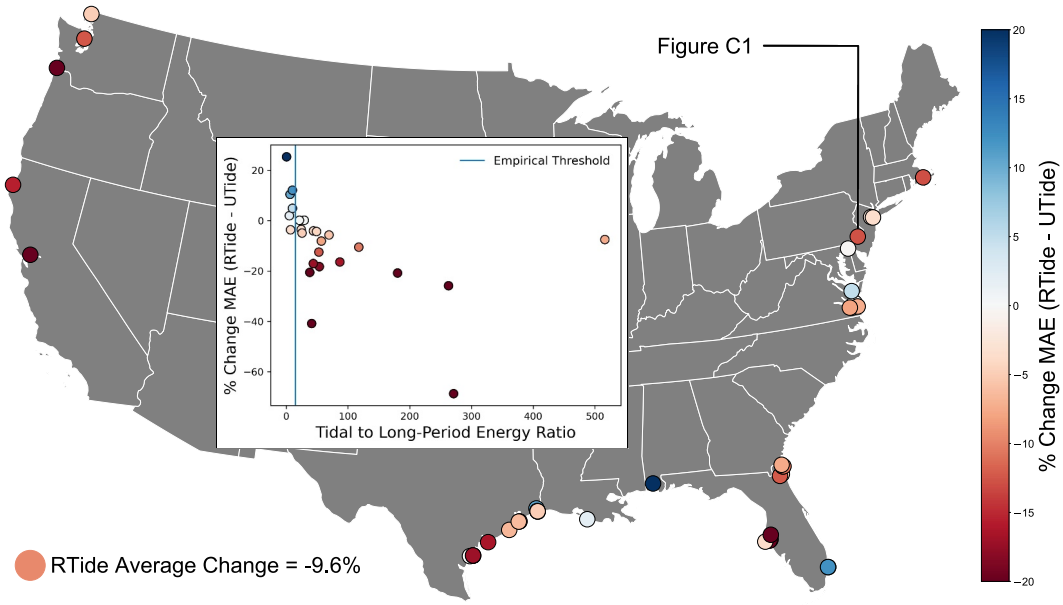

Figure 7. Comparison of RTide and UTide prediction error on 40 active NOAA current stations. Gauges are shaded according to the relative percent change in MAE between RTide and UTide. Red indicates a lower RTide MAE, and blue a higher MAE for example, UTide is better. A single MAE value is estimated by taking the Euclidean norm of the MAE of the $\hat{u}$ and $\hat{v}$ predictions. The inset plot shows the same station percentage change of MAE as a function of the ratio between Tidal and Long period energy as computed by Lomb-Scargle Periodogram.

spectral energy over two bands: a tidal band (0.8–12 cycles per day) and a low-frequency band (0.001–0.1 cycles per day), before calculating their ratio and averaging the results across both components. Partitioning the stations in this way produces a clear tidal-low-frequency threshold at approximately 15 in which the RTide approach yields consistently better predictive accuracy. Beneath this threshold (accounting for fewer than 20% of the given stations) the UTide models generally outperform RTide. This is in part due to the stability brought by the harmonic basis functions. It is worth noting that neither approach is effective at these stations as non-tidal processes are more dominant. In this context, we are utilizing the ratio of power spectral density as a proxy for how non-stationary the signal at a given station is. If the record at a given station is uniformly sampled, then a wavelet approach would be superior for this task (Lobo et al., 2024). Regardless, we advocate the usage of this simple check to evaluate the appropriateness of applying the astronomical tide only approach. If the tidal-long-period ratio is below 15, one should either consider the established nonstationary tidal analysis approaches (Matte et al., 2013), or consider including such forcing directly into RTide as described below. We provide a brief example of how nonstationary forcing can be built into RTide for sites located in tidal rivers in Appendix C.

5. Discussion

The preceding case studies have evaluated the capabilities of the response method for the prediction of tidal currents. The results offer a nuanced view of where and how the approach is best utilized. For fast-moving and strongly non-linear tidal currents, our response approach generally yields superior performance to HA in terms of mean absolute error. Additionally, for short reference series, the coupled response model demonstrates significant improvements in predictive performance over classical HA and in the accuracy and robustness of the derived constituents. For long reference series at less complex sites, the improvements are more moderate, and in a few cases, worse. In the case where nontidal forcing has a dominant effect, the gravitational approach described is not appropriate. Prior work has described the difference between harmonic and response methods as being akin to the difference between a Kodak (Harmonic) and Hasselblad (Response) camera (Zetler et al., 1979). Such a comparison seems appropriate for the classical method which required expert input to yield satisfactory response-based analysis. However, for all case studies in this manuscript, the same default set of parameters has been used. In this vein, we argue that our non-parametric, data-driven approach—implemented in RTide—offers the

same "point-and-shoot" capabilities as the harmonic method, enabling non-experts to benefit from the response approach, which is key for use in industry.

5.1. When and When Not to Use RTide

It is worth considering how the RTide program fits into the existing suite of tidal analysis programs and the conditions where it will and will not be successful. RTide is primarily designed for prediction and can be applied to non-uniformly sampled data. While there are unique physical-insights afforded by the learned response models, the non-parametric models can only be studied *ex post facto* in contrast to classical approaches which provide users the opportunity to specify a model a priori. Hence, if a user wishes to control which interactions are considered, the RTide approach should not be used. For engineering and operational contexts where maximizing the predictable variance is paramount, the RTide approach can significantly improve on conventional HA, as shown in this manuscript. Even when one is interested in just harmonic constituents, RTide can generally produce more accurate constituents. Furthermore, if one does not know the nature of the nonlinear interactions a priori, RTide offers a principled way to overcome this.

This manuscript does not consider the impact of non-tidal forcing which can significantly modulate the stationary tidal signal (Matte et al., 2013; Pan et al., 2018). Tidal rivers, meteorological effects and changing stratification are a few key examples. Under such conditions, the assumption of a time-invariant response to gravitational forcing is not valid. The pure tidal RTide approach (just using gravitational forcing) should not be used under these conditions. Here, one can either favor existing nonstationary approaches such as NS_Tide (Matte et al., 2013), or include non-stationary forcing as additional inputs into RTide. A basic example of this is given in Appendix C. Other approaches exist for studying non-stationary forcing such as wavelet-based approaches (Lobo et al., 2024), and even spline-based methods (Pan et al., 2018). While these approaches have limited utility for prediction, they are useful for quantifying whether the tidal signal is stationary and could be utilized to perform pilot analyses to decide if RTide is appropriate.

5.2. Further Considerations

Beyond predictive performance, there are notable trade-offs which must be acknowledged. First, the computational intensity of our approach is relatively higher. This is composed of three distinct components. The greatest expense is actually the computation of the input functions themselves. This expense is a weakness of our present implementation which recomputes the astronomical positions of the Sun and Moon at each individual time-step for each new location. Presently, the time expenditure for this step is on the order of minutes to an hour (in the case of a reference series containing several million measurements). Future work, will instead have these values precomputed which should significantly reduce this cost. Next, is the training of the model, which, even for extremely long reference series (millions of measurements), is on the order of seconds to minutes. Finally, predictions are near instantaneous and generally comparable to standard HA. In total, with the present implementation, this typically amounts to less than 1-min to completely run an RTide model on a single 30 days current reference series (3,600 measurements) using a single Apple M3 CPU. Figure S1 in Supporting Information S1, shows how this scales as a function of the number of observations. Due to this computational expense, for routine analysis tasks across many sites with sufficient data, HA is the preferred option. However, as shown and discussed, there are many specialist applications for which this increased computational cost is tolerable due to the performance improvements garnered.

Second, the learned response models, in the form of neural networks, are not easily interpreted. We have developed several techniques to get around this. The first, and perhaps the simplest, is to transform the response model into its harmonic equivalent. This is easily achieved by running HA on the response predictions. As noted above, this enables the equivalent HA model to inherit the improved frequency resolution characteristics of our coupled response approach. Additionally, as described in Monahan et al. (2025b), one can also apply *post hoc* model explainer tools such as SHAP (Lundberg, 2017). These tools can exploit the ability of a response approach to separate out the contributions of different forcing. This offers the potential to explicitly separate and study the response of tidal currents to non-gravitational forcing if it is included in the analysis.

With regard to non-tidal forcing, the response approach offers distinct advantages (Cartwright, 1968). The fundamental assumption of the response approach is that the oceanic response to forcing is time-invariant. If one only considers gravitational forcing, then this assumption can be invalidated if non-tidal forcing exists. However,

simple inspection of the shallow-water equations confirms this invariance holds for a wide range of compound tidal processes. Hence, in a response scheme, one must simply include the time-varying process as additional forcing to learn a coupled time-invariant response function. While the user must define an appropriate set of lags for these inputs, this can be achieved through an initial cross-spectral analysis. As shown in Monahan et al. (2025a), the non-parametric RTide approach makes this relatively simple as the time-series of the forcing can be directly included as an additional input into the model which was done for both fluvial and meteorological effects. An example of how this is applied to fluvial effects is given in Appendix C. The nature of the interactions of this forcing with the astronomical tide is determined from the data, rather than laborious analytical development. The direct incorporation of the tidal potential also removes the need for data from "uncontaminated" reference stations as in some methods. Forecasting, naturally relies on the ability to predict time-series of the corresponding forcing. Prior work for water levels has shown this is readily accomplished for tidal rivers and storm surge using the large suite of available hydrological and meteorological forecasts. A trade-off of this approach is that it loses the frequency interpretation of the model through classic harmonic analysis or the explicit definition of a functional form of the non-tidal interactions. These features are both implicit within NS_Tide (Matte et al., 2013). Given these trade-offs, the choice of which scheme to use rests on which features are useful to a given user and whether sufficient understanding exists to develop an analytical model for NS_Tide.

Due to the cubic relationship between current velocity and the power produced by tidal stream turbines, accurate tidal resource characterization is critical both for accurate prediction of power production and also the load on the turbine. As demonstrated, the coupled response models were able to replicate the performance in terms of mean absolute error of HA using $1/6$ of the data. In addition to tidal energy resource assessment, the ability to use sparse and short reference series may also be useful in deriving tidal currents from satellite observations (Hart-Davis et al., 2025). Separate work using this approach has also shown that similar improvements are achieved over operational numerical models (Monahan et al., 2025a). Such a result depends on the calibration of the numerical model used and cannot be seen as a general result. Future work is warranted to assess how a response approach can be combined with existing numerical models. One approach could be to replace the forcing from a small number of harmonic constituents at the boundary with a continuous response prediction. Consideration should also be given to post-processing of the numerical forecast which has been effective in global tide modeling (Hart-Davis et al., 2021). The ability to reduce the deployment duration of ADCPs for tidal energy resource assessment can lower the cost associated with the development phase of these sites and improve efficiency through better turbine placement. Future work will look to use the response method to quantify the impact of operational turbines on the tidal currents themselves. The quantification of the flow modulation from surges may also prove useful for protecting and maintaining deployed assets.

A final consideration when using RTide is the lack of uncertainty quantification. While the variational Bayesian HA procedure for obtaining harmonic equivalent models somewhat mitigates this, these estimates do not quantify the uncertainty in the response function itself. To overcome this several approaches can be tried. First, without modification, one can simply ensemble several RTide models with random initialization (Barber & Bishop, 1998). While this incurs greater computational cost, it provides a rigorous lower bound on the true marginal likelihood. An alternate approach is to explicitly model the time-invariant response as a stochastic process rather than a deterministic one. Experiments, not included in this work, have shown great promise in exploiting a class of methods termed *neural processes* which can approximate arbitrary stochastic processes within the RTide framework. Future work should evaluate these rigorously.

6. Conclusion

This manuscript presents a novel coupled model of tidal currents within the classic response method. Using our non-parametric and data-driven approach to weight estimation, the coupled response predictions demonstrate superior predictive accuracy over classical harmonic analysis and separable response models when predicting fast-moving tidal currents, whilst requiring considerably less data. Evaluating the pseudo-admittances of the coupled response model shows greater robustness when using short data series which is a consequence of the support of the tidal potential. We further show that the non-parametric weight estimation procedure advocated can effectively account for strong nonlinearities in the admittance function. Using ADCP data from the largest tidal energy site in the world, we show the coupled response model requires fewer than 30 days of measurements to achieve superior accuracy to HA trained on more than 180 days. The benefits of the response approach increase with the complexity of the currents at a given location. Our approach yields moderate improvements over HA for

operational forecasting tasks with an average reduction in absolute error of 9.6% on a sample of 40 active NOAA current stations. The modest increase in predictable variance captured is similar to but larger than that observed for tidal heights. A simple approach is also identified for determining whether a response approach will be effective for a given location. The brief case-study on applying our approach in tidal rivers highlights its potential for application in these regions. All response models can be implemented using just three lines of code in the open-source Python package RTide.

Appendix A: Error Metrics

Outlined below are the formulas for the metrics used for model validation.

1. The mean absolute errors (MAE) for each velocity component are given by

$$\mathrm{MAE}(\vec{u}) = \frac{1}{N}\sum_{i=1}^{N}|u_i - \hat{u}_i| \tag{A1}$$

$$\mathrm{MAE}(\vec{v}) = \frac{1}{N}\sum_{i=1}^{N}|v_i - \hat{v}_i| \tag{A2}$$

where N is the number of observations, u_i, v_i are the observed velocities, and $\hat{u}_i$, $\hat{v}_i$ are the predicted velocities. The total velocity RMS error is given by the Euclidean norm such that

$$MAE_E = \sqrt{MAE(\hat{u})^2 + MAE(\hat{v})^2} \tag{A3}$$

2. For each constituent k, the RMS error is computed independently for the orthogonal velocities $\vec{u}$ and $\vec{v}$ with

$$\mathrm{RMS}_k(\vec{u}) = \sqrt{\left[\left(A_u^{\mathrm{model}}\sin(\omega_k t) + B_u^{\mathrm{model}}\cos(\omega_k t)\right) - \left(A_u^{\mathrm{true}}\sin(\omega_k t) + B_u^{\mathrm{true}}\cos(\omega_k t)\right)\right]^2} \tag{A4}$$

$$\mathrm{RMS}_k(\vec{v}) = \sqrt{\left[\left(A_v^{\mathrm{model}}\sin(\omega_k t) + B_v^{\mathrm{model}}\cos(\omega_k t)\right) - \left(A_v^{\mathrm{true}}\sin(\omega_k t) + B_v^{\mathrm{true}}\cos(\omega_k t)\right)\right]^2} \tag{A5}$$

where the overbar indicates an average over a full-period ($0 \rightarrow 2\pi$) for the given constituent. The total RMS error is given by the Euclidian norm of the two RMS values such that

$$\text{Total } RMS_k = \sqrt{RMS_k(\vec{u})^2 + RMS_k(\vec{v})^2} \tag{A6}$$

Appendix B: Admittances

As mentioned in the main text, it is worth evaluating whether the non-parametric response approach used in RTide can capture sharp deviations in the admittance function. Psuedo-Admittances computed from harmonic constituents estimated from simulated data from Site C in the Pentland simulation are shown in Figure B1. It can clearly be seen that the coupled RTide approach captures a majority of the variance associated the with MSN2 constituent despite it deviating from other values in the tidal potential by two orders of magnitude. A corresponding figure (Figure S2 in Supporting Information S1) is also given in the Supporting Information S1 which omits MSN2 and uses a simple polynomial curve fit to visualize the smooth portion of the admittance.

Another important consideration is how the response and harmonic approaches behave as the amount of data is decreased. In the low-data limit, least-squares estimators typically exhibit a failure mode where the lack of regularization can yield spurious constituent estimates which are extremely large (Monahan et al., 2025c). We illustrate this behavior, and compare it to RTide, in Figures S4 and S5 of Supporting Information S1.

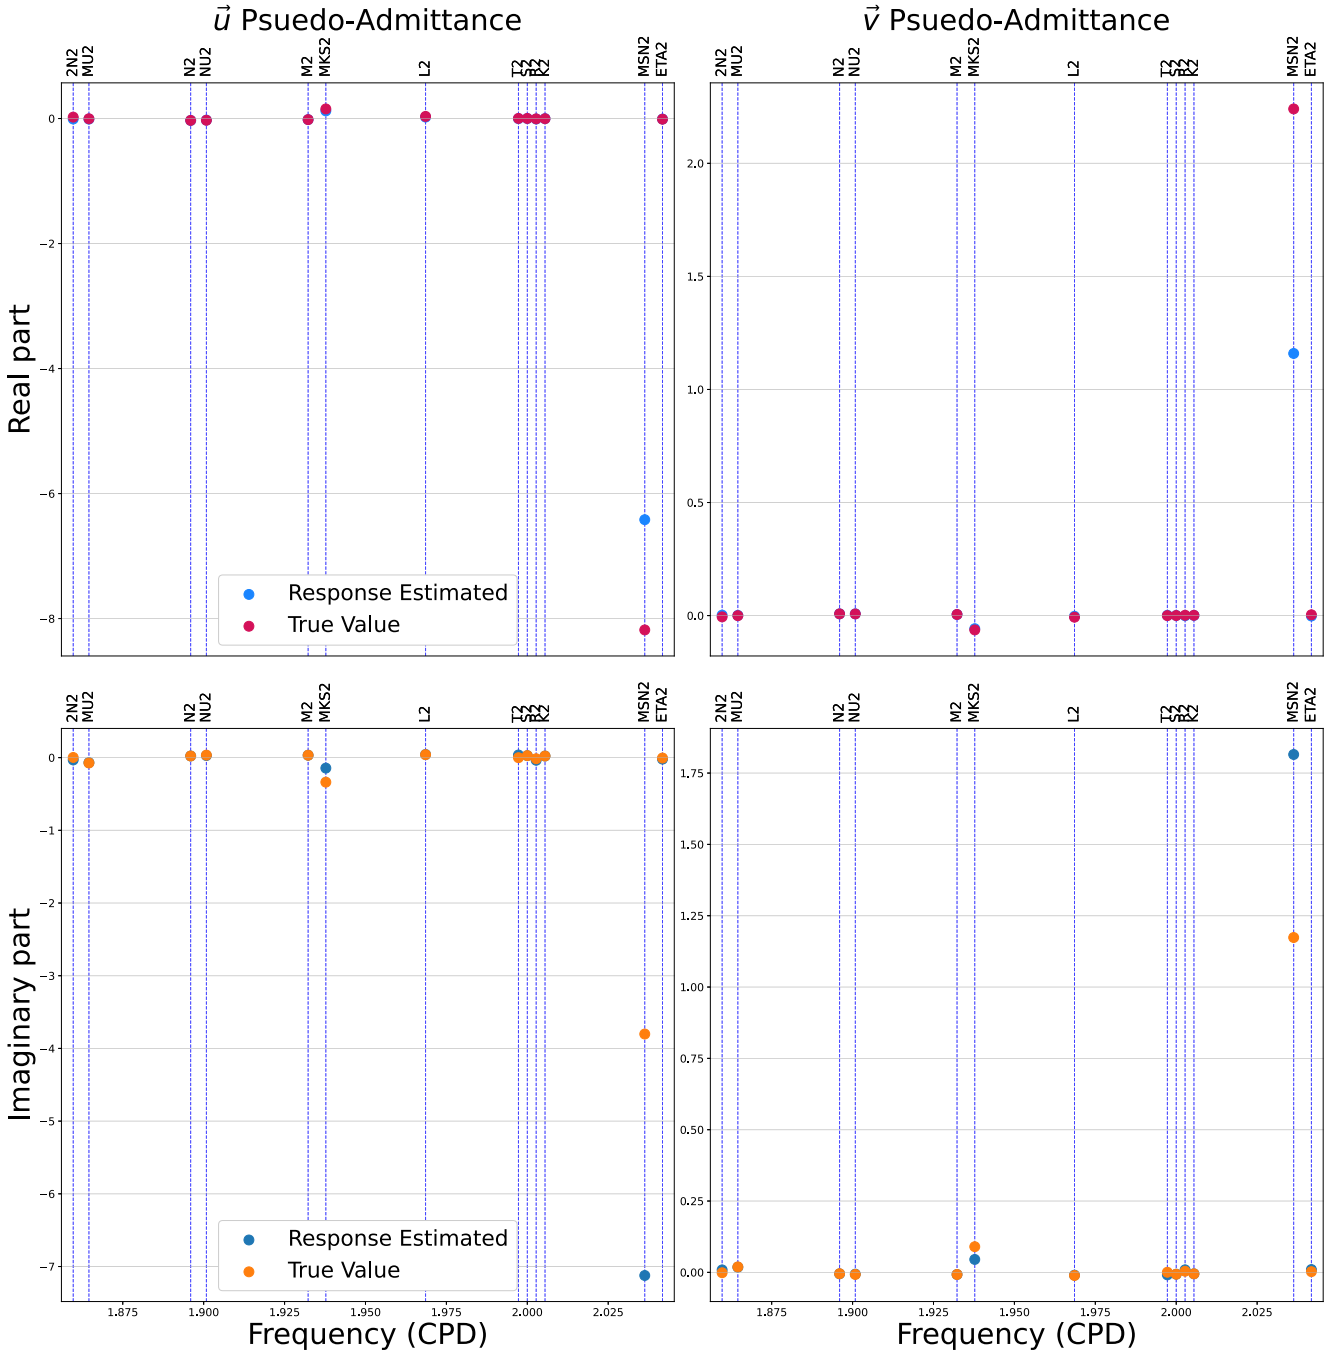

Figure B1. Psuedo-Admittance comparisons between standard coupled RTide (Response Estimated) and true harmonic constituents (True Value) at simulation location C. Psuedo-admittances are computed by dividing the quadrature amplitudes of each constituent and each velocity component by its magnitude in the tidal potential (taken from (D. Cartwright & Edden, 1973)). Note many Response estimates are hidden by the true value due to the good agreement.

Interestingly, we show that the RTide approach yields more robust Psuedo-admittances even using 1-day of observations with added noise within the standard ADCP error range.

Appendix C: Tidal Rivers

We here consider a brief case-study of how nonstationary forcing can be included in the response approach to support our claims in the discussion. To evaluate this we focus on the tidal currents in Philadelphia, PA, USA.

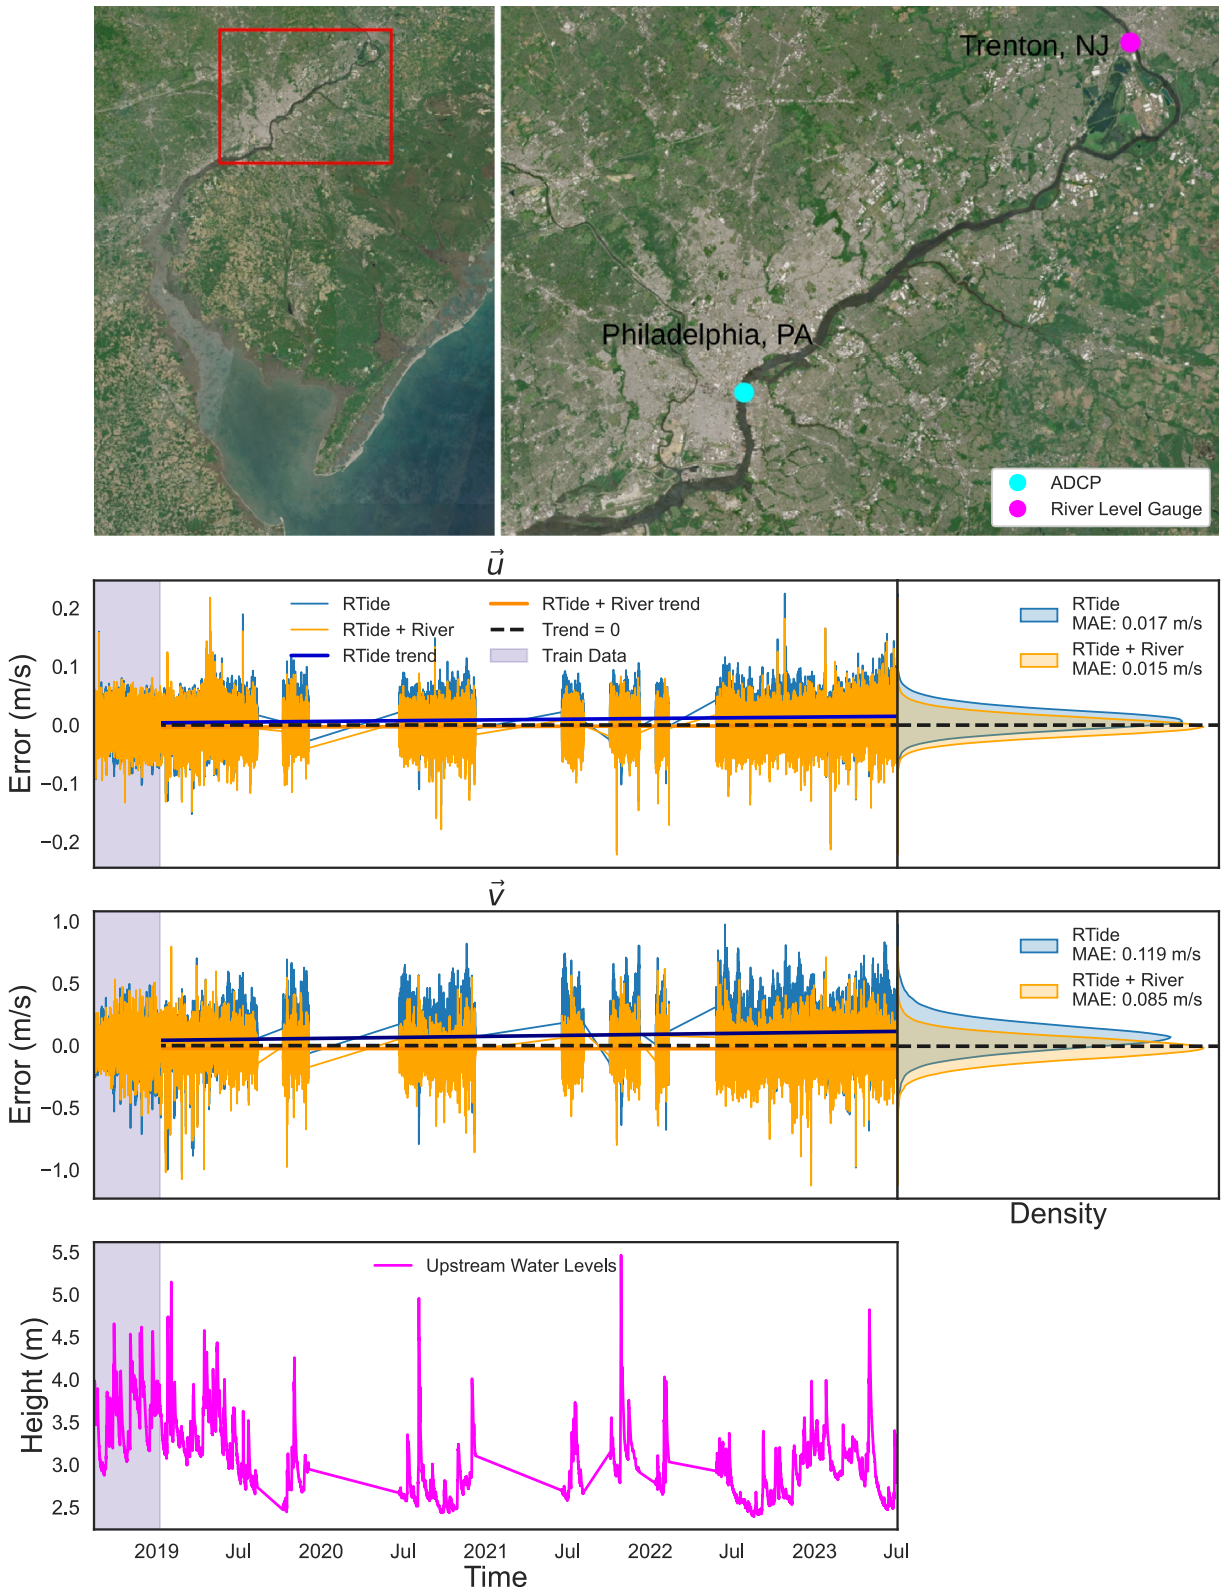

Figure C1. Evaluation of RTide in a tidal river using upstream river water levels as forcing. Top row: Images of the study region with location of ADCP and upstream river level gauge overlaid from Sentinel-2 (Phiri et al., 2020). Middle panels: Time-series of model residuals. RTide uses only gravitational forcing. RTide + River receives the instantaneous upstream river water levels as an additional input. Both models are trained on 120 days of data shown in the shaded region. Trends and KDE estimates of the residuals only over the testing data are shown. Bottom Panel: Upstream river water levels.

This site was primarily selected for the fact that the ADCP (gauge id: db0301) and River level gauge (gauge id: USGS-01463500) data were easily accessible from NOAA and USGS online services respectively. However, the approach can be applied wherever concurrent upstream river level or discharge data is available (or conceivably other relevant types of forcing such as meteorological data).

Philadelphia is located far up the Delaware Bay estuary, which is the outlet of the Delaware river as is seen in the top panel of Figure C1. The tides in this region are non-stationary due to the varying river outflow. A visualization of a representative week of measurements is given in Figure S3 of Supporting Information S1. It is worth noting that the upstream water levels in Trenton exhibit minimal tidal influence. The non-stationarity introduced by the river can be accounted for in the response framework by learning the time-invariant and coupled response between the gravitational forcing and the river outflow. To do this, we simply include the real-time upstream water level in Trenton, NJ as an additional forcing input. The result of making this simple modification is seen in Figure C1. Here, RTide models are trained for 120 days on ADCP data from Philadelphia. The RTide + River model receives as additional input the upstream water levels at Trenton. Inspection of the residuals show that the addition of river water levels upstream reduced the out-of-sample predictive MAE by 28.2%. An encouraging result is that while the RTide + River model was trained on upstream water levels greater than 3 m over only 120 days, the residuals remain centered about zero across the full 4-year test period which was characterized by significant periods of upstream water levels below 3 m. Interestingly, this was a site for which RTide already considerably outperformed the corresponding HA. In this context, while the river outflow does contribute, it does not dominate the tidal signal as it does in other regions. Hence, future work is needed to evaluate how this simple approach to capturing these nonstationary effects can be applied in strongly nonstationary regions. The incorporation of upstream river levels also renders the harmonic equivalent model devised in this paper inaccurate. Future work should instead look to combine this with the NS_Tide framework to provide a principled way of interpreting the nonstationary RTide models.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

All code needed to replicate the results is included in the dedicated Zenodo repository https://doi.org/10.5281/zenodo.17477856. Additionally, the method is integrated into the RTide Python package (https://github.com/thomasmonahan/RTide). All NOAA gauge data can be accessed freely via the NOAA Tides and Currents website: https://tidesandcurrents.noaa.gov/. Simulation data from the Pentland Firth is available from the provided repository. The data from the MeyGen section of this study are available from SAE Renewables but restrictions apply to the availability of these data, which were used under academic partnership for the current study.

Acknowledgments
TT and TM acknowledge funding from the Schmidt AI in Science Fellowship, a Schmidt Futures program. TA was supported by the EPSRC CoTide project (EP/X03903X/1). We thank Fraser Johnson for sharing the data from the MeyGen site.

References

Adcock, T. A. A., Draper, S., Houlsby, G. T., Borthwick, A. G. L., & Serhadlıoğlu, S. (2013). The available power from tidal stream turbines in the Pentland firth. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *469*(2157), 20130072. https://doi.org/10.1098/rspa.2013.0072

Adcock, T. A. A., Draper, S., Houlsby, G. T., Borthwick, A. G. L., & Serhadlıoğlu, S. (2014). Tidal stream power in the Pentland firth–long-term variability, multiple constituents and capacity factor. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, *228*(8), 854–861. https://doi.org/10.1177/0957650914544347

Barber, D., & Bishop, C. M. (1998). Ensemble learning in Bayesian neural networks. *Nato ASI Series F Computer and Systems Sciences*, *168*, 215–238.

Blunden, L., & Bahaj, A. (2007). Tidal energy resource assessment for tidal stream generators. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, *221*(2), 137–146. https://doi.org/10.1243/09576509jpe332

Cartwright, D. (1982). Tidal analysis-A retrospect. In *Developments in water science* (Vol. 17, pp. 170–188). Elsevier. https://doi.org/10.1016/s0167-5648(08)70709-3

Cartwright, D., & Edden, A. C. (1973). Corrected tables of tidal harmonics. *Geophysical Journal International*, *33*(3), 253–264. https://doi.org/10.1111/j.1365-246x.1973.tb03420.x

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*, *95*(C3), 3069–3090. https://doi.org/10.1029/jc095ic03p03069

Codiga, D. L. (2011). Unified tidal analysis and prediction using the UTide Matlab functions. https://www.mathworks.com/matlabcentral/fileexchange/46523-utide-unified-tidal-analysis-and-prediction-functions

Flinchem, E., & Jay, D. (2000). An introduction to wavelet transform tidal analysis methods. *Estuarine, Coastal and Shelf Science*, *51*(2), 177–200. https://doi.org/10.1006/ecss.2000.0586

Godin, G. (1991). The analysis of tides and currents. *Tidal hydrodynamics*, *675*, 709.

Hart-Davis, M. G., Moiseev, A., Bonaduce, A., Backeberg, B. C., & Johannessen, J. A. (2025). Direct ocean tidal current measurements from space: Enhanced interpretation of sentinel-1 Doppler shift signals. *Geophysical Research Letters*, *52*(13), e2025GL115779. https://doi.org/10.1029/2025gl115779

Hart-Davis, M. G., Piccioni, G., Dettmering, D., Schwatke, C., Passaro, M., & Seitz, F. (2021). Eot20: A global ocean tide model from multi-mission satellite altimetry. *Earth System Science Data*, *13*(8), 3869–3884. https://doi.org/10.5194/essd-13-3869-2021

Innocenti, S., Matte, P., Fortin, V., & Bernier, N. (2022). Analytical and residual bootstrap methods for parameter uncertainty assessment in tidal analysis with temporally correlated noise. *Journal of Atmospheric and Oceanic Technology*, *39*(10), 1457–1481. https://doi.org/10.1175/jtech-d-21-0060.1

Jay, D. A., & Flinchem, E. P. (1997). Interaction of fluctuating river flow with a barotropic tide: A demonstration of wavelet tidal analysis methods. *Journal of Geophysical Research*, *102*(C3), 5705–5720. https://doi.org/10.1029/96jc00496

Jay, D. A., & Musiak, J. D. (1996). Internal tidal asymmetry in channel flows: Origins and consequences. *Mixing in estuaries and coastal seas*, *50*, 211–249. https://doi.org/10.1029/ce050p0211

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795. https://doi.org/10.2307/2291091

Klema, M. R., Pirzado, A. G., Venayagamoorthy, S. K., & Gates, T. K. (2020). Analysis of acoustic Doppler current profiler mean velocity measurements in shallow flows. *Flow Measurement and Instrumentation*, *74*, 101755. https://doi.org/10.1016/j.flowmeasinst.2020.101755

Leffler, K. E., & Jay, D. A. (2009). Enhancing tidal harmonic analysis: Robust (hybrid L1/L2) solutions. *Continental Shelf Research*, *29*(1), 78–88.

Le Provost, C. (1991). And compound tides. *Tidal hydrodynamics*, *269*.

Le Provost, C., Lyard, F., & Molines, J.-M. (1991). Improving ocean tide predictions by using additional semidiurnal constituents from spline interpolation in the frequency domain. *Geophysical Research Letters*, *18*(5), 845–848. https://doi.org/10.1029/91gl01065

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

Lundberg, S. (2017). A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*. https://doi.org/10.48550/arXiv.1705.07874

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., Jay, D. A., & Zaron, E. D. (2013). Adaptation of classical tidal harmonic analysis to nonstationary tides, with application to river tides. *Journal of Atmospheric and Oceanic Technology*, *30*(3), 569–589. https://doi.org/10.1175/jtech-d-12-00016.1

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. A. (2025a). Prediction of tidal currents in the Inner sound of the Pentland firth using RTide. In *Proceedings of the 16th European wave and Tidal Energy Conference, funchal*.

Monahan, T., Tang, T., Roberts, S., & Adcock, T. A. A. (2025b). RTide: Automating the tidal response method. *Journal of Geophysical Research: Machine Learning and Computation*, *2*(2), e2024JH000525. https://doi.org/10.1029/2024jh000525

Monahan, T., Tang, T., Roberts, S., & Adcock, T. A. A. (2025c). 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., & Hasselmann, K. (1964). Super-resolution of tides. *Studies on oceanography*, 339–344.

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.

Phiri, D., Simwanda, M., Salekin, S., Nyirenda, V. R., Murayama, Y., & Ranagalage, M. (2020). Sentinel-2 data for land cover/use mapping: A review. *Remote Sensing*, *12*(14), 2291. https://doi.org/10.3390/rs12142291

Prandle, D. (1982). The vertical structure of tidal currents. *Geophysical & Astrophysical Fluid Dynamics*, *22*(1–2), 29–49. https://doi.org/10.1080/03091928208221735

Prandle, D. (1997). Tidal currents in shelf seas—their nature and impacts. *Progress in Oceanography*, *40*(1–4), 245–261. https://doi.org/10.1016/s0079-6611(98)00013-5

Prandle, D., & Matthews, J. (1990). The dynamics of nearshore surface currents generated by tides, wind and horizontal density gradients. *Continental Shelf Research*, *10*(7), 665–681. https://doi.org/10.1016/0278-4343(90)90044-m

Rhodes, B. (2019). Skyfield: High precision research-grade positions for planets and Earth satellites generator. *Astrophysics Source Code Library, ASCL–1907*. https://ui.adsabs.harvard.edu/abs/2019ascl.soft07024R/abstract

Roberts, S. J., & Penny, W. D. (2002). Variational Bayes for generalized autoregressive models. *IEEE Transactions on Signal Processing*, *50*(9), 2245–2257. https://doi.org/10.1109/tsp.2002.801921

Schendel, A., Hildebrandt, A., Goseberg, N., & Schlurmann, T. (2018). Processes and evolution of scour around a monopile induced by tidal currents. *Coastal Engineering*, *139*, 65–84. https://doi.org/10.1016/j.coastaleng.2018.05.004

Sitzmann, V., Martel, J., Bergman, A., Lindell, D., & Wetzstein, G. (2020). Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, *33*, 7462–7473.

Stiven, T., Couch, S. J., & Iyer, A. S. (2011). Assessing the impact of ADCP resolution and sampling rate on tidal current energy project economics. In *OCEANS 2011 IEEE-Spain* (pp. 1–10).

Thomson, C. (2016). Turn of the tide. *Engineering & Technology*, *11*(11), 50–53.

Woodworth, P. L. (2012). A note on the nodal tide in sea level records. *Journal of Coastal Research*, *28*(2), 316–323. https://doi.org/10.2112/jcoastres-d-11a-00023.1

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

Zetler, B., Cartwright, D., & Berkman, S. (1979). Some comparisons of response and harmonic tide predictions. *International Hydrographic Review*.

Zetler, B. D., & Munk, W. H. (1975). The optimum wiggliness of tidal admittances. *Journal of Marine Research*, *33*.