# Spatiotemporal tidal prediction and analysis through physics-informed ML

Thomas Monahan

St. Cross College

University of Oxford

Supervised by Thomas Adcock, Stephen Roberts, and Tianning Tang

A thesis submitted for the degree of Doctor of Philosophy

Hilary Term, 2025

# Acknowledgements

This thesis would not have been possible without the mountain of support provided by my family, friends, colleagues, and supervisors.

First, this thesis, and indeed all that I have accomplished has been built on the love and support from my family. Sonia, Mike and John, thank you for empowering me to go for my goals and for always being there when they falter.

Second, I would like to thank all of the friends I have been fortunate to cross-paths with during this DPhil. The civil engineering lunch group, and fun lunch Fridays, is truly a highlight of each day – particularly the post pub trips to Pepper's Burgers. In particular, I should like to thank Dan, Rachel, and Tim whose crumpet breaks and random chats made even the dreariest of British days enjoyable. Thank you to all of my teammates on Wolfson College Boat Club, and in particular David, Dom, Gio, Adam, and Ed for the much needed distraction and putting up with my bad mood in the mornings. To the friends I escaped the rain in Oxford with; Luca, Brennan, Steven, Xavier, Kevin and my conference buddy Mike, thank you and I look forward to many more escapes. Next, I wish to thank Nadav, whose unwavering support and friendship, even from a far, has been a constant in my life. Finally, to Mary, who has been with me nearly every step of the way through this DPhil – it has meant the world to share this journey with you. Your love and support has gotten me through the highs and lows of these 2.5 years. Whatever the future may have in store for us, I am grateful to have you in my corner.

I must extend my gratitude to my three wonderful supervisors Thomas Adcock, Tim (Tianning) Tang and Stephen Roberts. Tom, you took a shot on me when I on-paper did not deserve it. I am grateful for your support, mentorship, and for championing me in my various endeavors. Tim, it has been a privilege to be your first student – I and I'm certain many future students are/will be grateful for your kindness and insightful ideas. Steve, thank you for the always fascinating meetings, for pushing me to develop and understand these methods better, and for somehow always finding time to meet. I greatly look forward to working with all of you in the future. I should also like to thank Ronald Roy, who through a serendipitous connection to Lewiston, ME, became a mentor of mine. Ron, thank you for your mentorship, and delightful coffee chats about Maine. I hope retirement treats you well.

I would like to thank the Engineering Department for funding my DPhil. Further, I should like to thank Clive and Alison for their assistance, coffee chats, and the occasional angle grinder.

# Abstract

Tidal processes—those driven by gravitational forcing from the Earth, Moon, and Sun—shape coastal environments and significantly impact infrastructure. While conventional methods of tidal analysis and prediction perform well for barotropic tides sampled at high temporal resolution, major challenges arise when exogenous non-tidal forcing is present or temporal sampling becomes sparse and irregular. Under these conditions, the accuracy of empirical methods degrades as their underlying assumptions are violated.

This thesis introduces four novel methodologies that advance the prediction and analysis of tidal and tidally driven processes across diverse spatiotemporal scales. Central to each is the augmentation of classical empirical techniques with modern machine learning, reducing reliance on inaccurate assumptions and improving both predictive power and physical interpretability.

The work focuses on two classical approaches to tidal analysis: harmonic analysis and the response method. Each has strengths that suit particular challenges. For harmonic analysis, I develop a Bayesian framework tailored to sparse, noisy satellite altimetry data. This approach enables the exploitation of Surface Water and Ocean Topography (SWOT) mission data, provides uncertainty estimates lacking in existing models, and led to the discovery of an earth-shaking wave using this dataset.

For the response method, I develop an automated, non-parametric procedure—overcoming a major limitation of the classical approach. It is the only method capable of analyzing and predicting arbitrary tidal processes, including tidal rivers and storm surges. I further propose a coupled response theory for predicting tidal currents. This method offers superior predictive accuracy to harmonic analysis and reduces the required time-series duration by nearly an order of magnitude – a feature that can enhance harmonic analysis as well.

Finally, I consider augmenting all classical methods – including harmonic, response, and numerical models – with short-term forecasting using recent data via an online training procedure. These contributions address long-standing limitations in empirical tidal analysis and have broad applications – from improving satellite data usage to enhancing operational storm surge forecasting. This thesis argues that the 'tide problem' remains unsolved and offers tools and solutions to some of its most pressing challenges.

# Contents

# List of Figures

# List of Tables

# 1 Introduction

Tidal processes shape our coastlines, influence global climate systems, and drive critical oceanic processes—yet despite centuries of study, accurately predicting their behavior using either numerical or empirical approaches remains a fundamental challenge. While the methods available for studying these phenomena have remained relatively the same over the past decades, countless new observational modalities exist, with new satellite missions producing exponentially more data than ever before. This thesis develops three new methods for the prediction and analysis of tides and tidally driven processes across a range of spatiotemporal scales. The applications of these methods are broad and far reaching, ranging from improving the usage of satellite altimetry data to operational storm surge prediction. Central to this work is the augmentation of classical methods of empirical analysis using modern data-driven machine learning. As will be shown, usage of these methods actually enables more of the "physics" to be kept in the problem by (i) reducing the limiting assumptions made by the model, and (ii) allowing flexibility with which to explicitly model departures from theory. These advantages not only yield superior predictive accuracy, but offer several new avenues for analysis of the learned model dynamics which have not previously been possible. Here, I briefly introduce the three classic approaches to analyzing and predicting tides as well as their primary limitations. The theory for each approach will be provided in full in the literature review chapter, as well as in the respective chapters.

Tides are widely considered to be the most predictable of all geophysical phenomenon (10). This is a consequence of the fact that tides are produced by the combined gravitational forcing from the periodic movements of the earth, moon, and sun – which are fundamentally predictable. In most regions globally, it is reasonable (and often necessary) to assume the tidal response to this forcing is coherent with the forcing itself. That is, tides will exist at the same discrete frequencies as the astronomical forcing. This gives rise to several natural, and effective analysis procedures: namely harmonic analysis (11). As the name suggests, harmonic analysis approximates tides to be the superposition of tidal waves (termed harmonic constituents) at these discrete frequencies. Nonlinearity in the tidal spectrum is handled by taking multiples, sums and differences, of the linear constituents. In total, the tidal spectrum can be divided into more than 600 frequencies, the resolution of which is fundamentally linked to the length of the tidal record (12). Modern harmonic analysis solves for these constituents in the time domain using least-squares estimators. The accuracy of the derived constituents degrades significantly in the presence of non-stationary noise and when applied to sparsely sampled data due to biasing of the estimator (13). Unfortunately, these conditions are also commonly encountered when analyzing satellite altimetry data. Chapter 3 develops a novel Bayesian approach to handling these data (13), and Chapter 4 applies it to discover a new phenomenon (14).

As noted above, the underlying assumption of harmonic analysis is that the tidal response is coherent with the astronomical forcing, and thus periodic. Unfortunately, due to interactions with other phenomena (e.g. river outflow, atmospheric forcing, etc.), this is not always the case. In fact, in these regions the tidal profile can become jagged and strongly non-stationary (See Figure 5.3) (15). Hence, conventional harmonic methods are no longer viable. In the 1960's two greats of oceanography, Munk and Cartwright, developed the response method (5). As the name suggests, the approach casts the tidal prediction problem within an impulse-response framework under the assumption that the oceanic response to such forcing is time-invariant. Here, the "impulse" is computed directly from the Kepler-Newtonian mechanics and is expanded in terms of spherical harmonics. The response prediction is thus defined as the convolution between past, present (and indeed sometimes future) values of the tidal potential. Weights are estimated in several ways using conventional spectral analysis techniques. This approach has several intrinsic advantages, namely the removal of the assumption that the oceanic response is coherent with the astronomical forcing. This makes it possible, in theory, to apply the response method to these non-stationary problems. However, non-stationary forcing also introduces nonlinearity which is difficult to represent parametrically in the conventional response formalism. Consequently, the response method has not been used for such applications beyond its initial case study (16). To combat this, I develop a nonparametric response approach in Chapter 5 ((17)). This same approach is extended to the prediction of tidal currents, in conjunction with a novel "coupled" response theory of tidal currents presented in Chapter 6.

It is also possible to predict tides and tidally driven processes using numerical modeling approaches. Here, I use "numerical modeling" as a term to encompass both fully numerical and assimilative approaches. Indeed, numerical modeling is the basis of almost all modern storm surge forecast systems (18). Here, the oceanic response is explicitly computed from both the gravitational (tidal) and atmospheric forcing using either 2-d or 3-d equations governing tide-surge motion – typically the shallow water equations. While these procedures are physically grounded, they are limited by (i) the accuracy of the boundary conditions employed, especially bathymetry data, (ii) the resolution of the mesh used, and (iii) the approximations used in deriving the flow equations. The boundary conditions are particularly problematic when developing global models as the quality of bathymetric data varies massively. The resolution issue is perhaps the most significant limitation of these approaches as computational bottlenecks necessitate the usage of coarse grids. Finally, approximations in the flow equations can lead to the neglect of physical processes altogether. As noted above, the complexity of tides in coastal regions leads to significant modeling errors and the superiority of empirical approaches in these regions (19). In fact, modern surge models do not actually use the tide component computed by their numerical models due to the poor accuracy, instead adding the computed surge onto harmonic predictions computed from gauges (18). This problem is explicitly treated in Chapter 5 (17), and a method which can greatly augment the performance of conventional numerical surge forecasts across any forecast horizon is produced. It should be noted that similar numerical models are also utilized for forecasting tidal currents. Short-term tidal current forecasts are relevant for tidal energy site operations and load forecasting. As such, in addition to the aforementioned limitations, numerical models are also unable to adapt to these variable conditions. A specialized procedure for this application is introduced in Chapter 7 (20).

From the outset of this thesis, it is worth addressing the state of modern tidal analysis, and why this work is needed. The 'tide problem' has a long history of being 'solved'. The first closing of the tide problem followed the publication of Newton's Principia in 1687. The problem was then reopened, and again, solved in Pierre-Simon Laplace's 1799 Mécanique Céleste. Walter Munk noted that he was "...told by US funding agencies in the 1960s that the tidal problem had been solved by Victorian mathematicians..." (21). Notably, this statement just predates his and Cartwright's breakthrough development of the response method (5). From this development, and the advent of satellite altimetry, came a renaissance period in tidal analysis. In the subsequent decades, our understanding (and ability to predict) global tides has advanced massively. However, many problems remain in coastal and estuarine regions, where nonlinearity and non-tidal forcing strongly impact tidal propagation, degrading the efficacy of conventional methods.

These deficiencies create problems across a multitude of applications. My own work (using the method developed in Chapter 5) improving storm surge forecast accuracy at the Thames Barrier in London exemplifies this, and the significant implications these deficiencies have for public safety. Additionally, with the proliferation of marine renewable energy in recent years, the need to understand the distribution of the tidal resource in complex regions has never been greater. This is critical not just for extraction, but to mitigate the impact of tides on off-shore infrastructure. Lastly, while new satellite altimetry missions provide unprecedented views into coastal dynamics, the deficiency in coastal tide modeling limits our ability to fully exploit these data. Hence, this thesis argues that the ‘tide problem’ is not solved. To be regarded as such, adequate methods need to exist for the analysis and prediction of these more complex tidal and tidally driven processes which have significant implications for communities and industry.

## 1.1 Thesis outline

This integrated thesis is structured into five chapters, alongside a literature review, and conclusion. The five central chapters are given by published or under-review journal articles. As discussed above, while the methods developed herein can enhance classical tidal analysis tasks, the objective of this work is to address the ‘messy’ problems in tidal prediction and analysis—where the assumptions of classical approaches are limiting, and the impacts on humans and industry are most significant. As an integrated thesis, each chapter represents a different paper, and is thus self-contained. This may lead to some inconsistency in notation between chapters. A description of the content of each chapter is as follows, for the content Chapters 3-7 these are abstracts.

### 1.1.1 Chapter 2: Literature Review

This chapter introduces the physical theory of tides and surges. This context provides a useful intuition for the two empirical approaches considered in this thesis – Harmonic Analysis and the Response method. Using this physical basis, I examine the physical assumptions made by each approach and outline the various methods which have been developed to overcome their intrinsic limitations. While this thesis considers applications for a broad range of phenomena and contexts, the motivation for each of this thesis' developments is easily seen from these comparisons. Fundamental to understanding the limitations of these empirical approaches is the consideration of the various mathematical machinery employed for analysis, which I also discuss. I also consider the relatively recent usage of machine learning and other approaches such as wavelet analysis. The strengths and weaknesses of each method are discussed, and the research gaps which will be addressed in subsequent chapters are clearly outlined.

### 1.1.2 Chapter 3: Tidal corrections from and for SWOT using a spatially coherent variational Bayesian harmonic analysis

The accuracy of global tidal models degrades significantly in coastal and estuarine regions. These models are important for correcting measurements from satellite altimetry and are used in numerous scientific and engineering applications. The new Surface Water Ocean Topography (SWOT) mission is providing measurements at unprecedented horizontal resolution in these regions. This dataset provides both the opportunity and the need to quantify and correct the spatial variability of these errors. I develop a variational Bayesian framework for tidal harmonic analysis which can be applied in the early stages of SWOT. The approach demonstrates superior robustness to different types of noise contamination in comparison to conventional least-squares approaches whilst providing full uncertainty estimation. By imposing a spatially coherent inductive bias on the model, I achieve superior harmonic constituent inference using extremely sparsely sampled data. Bayesian uncertainty estimation gives rise to statistical methods for outlier removal and constituent selection. Using our approach, I estimate a lower bound for the residual tidal variability for two SWOT Cal/Val passes (003 and 016) around the European Shelf to be at least 7% on average. I also show similar estimates cannot be produced using standard least-squares approaches. Tide gauge validation in the same region confirms the superiority of our empirical approach in coastal environments. Empirical corrections for the SWOT data products are provided alongside an open-source Python package, VTide[1].

### 1.1.3 Chapter 4: Observations of the seiche that shook the world

On September 16th, 2023, an anomalous 10.88 mHz seismic signal was observed globally, persisting for 9 days. One month later an identical signal appeared, lasting for another week.

---

[1] `https://github.com/thomasmonahan/VTide`

Several studies have theorized that these signals were produced by seiches which formed after two landslide-generated mega-tsunamis in an East Greenland fjord. This theory is supported by seismic inversions, and analytical and numerical modeling, but no direct observations have been made – until now. Using data from the new Surface Water Ocean Topography mission, I present the first observations of this phenomenon. By ruling out other oceanographic processes, I validate the seiche theory of previous authors and independently estimate its initial amplitude at 7.9 m using Bayesian machine learning and seismic data. This study demonstrates the value of satellite altimetry for studying extreme events, while also highlighting the need for specialized methods to address the altimetric data's limitations, namely temporal sparsity. These data and approaches will help in understanding future unseen extremes driven by climate change.

### 1.1.4 Chapter 5: Automating the tidal response method

Tides pose significant operational and engineering challenges and are critical drivers of many natural processes. Accurate tidal predictions are important for modeling these phenomena. Conventionally, tidal prediction is carried out using harmonic analysis, the accuracy of which degrades when non-stationary and nontidal forcing are present. While Munk and Cartwright's response method avoids the assumptions that give rise to this degradation, the difficulty of defining realistic interactions between inputs has inhibited automated applications. Here, I develop a non-parametric framework for tidal analysis and prediction of sea levels under compound forcing. The approach embeds a class of neural networks capable of representing any arbitrary Volterra series—the mathematical basis of Munk and Cartwright's response method—within the classic method. The new ML Response Framework overcomes the automation challenges imposed by the original method and allows for the direct inference of high-order nonlinear interactions. This approach makes the inclusion of meteorological and other non-tidal forcing straightforward. Furthermore, I show that by accounting for this variability explicitly, a better astronomical tidal estimate is obtained. A method is devised to obtain physical insights from the learned model, illustrating how it can be used to study the interaction and modulation of astronomical tides by external forcing. By removing the need to specify a parametric form of the regression beforehand, our approach makes the study of phenomena that heretofore could not be accounted for straightforward. I provide several case studies, including the analysis and prediction of tide-surge interaction, riverine tides, and nuisance flooding. These applications, and more, can be replicated using only three lines of code with the open-source Python package (RTide)[2].

### 1.1.5 Chapter 6: Response-based prediction of tidal currents

This study presents the first comprehensive evaluation of the response method for predicting tidal currents. I introduce a novel “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. This performance advantage increases with the complexity of the tidal currents, offering more modest reductions in absolute error of 9.6% on average across 40 active NOAA current stations. The framework offers several opportunities for future work in understanding the role of non-tidal forcing and sediment transport and has significant economic implications for tidal energy site development. The proposed approach is implemented and freely available in the open-source RTide Python package. Critically, the non-parametric approach removes the need for expertise when using the response method.

### 1.1.6 Chapter 7: A hybrid model for Online short-term tidal energy forecasting

A hybrid model is proposed for the short-term online prediction of tidal currents. The harmonic residual analysis (HRA) model is designed to augment the numerical schemes employed by tidal energy installations by forecasting the residual error of existing methods. Using a combination of techniques from Information and Fractal Theory, a novel component selection criterion for singular spectrum analysis (SSA) is used to remove true noise from the residual time series and to decompose the signal into components that are appropriate for linear-recurrent forecasting (LRF) and high order fuzzy time series (HOFTS) respectively. The performance of the HRA method is evaluated using a combination of simulated and real data from sites in the United Kingdom and the United States. Results demonstrate the model’s viability for 6-minute and 1-hour forecast horizons across sites exhibiting variable degrees of non-linearity.

---

[2] `https://github.com/thomasmonahan/RTide`

### 1.1.7 Chapter 8: Discussion and Conclusion

This chapter brings each method into discussion with the broader challenges in tidal analysis and prediction. I highlight the strengths and weaknesses of each approach, and offer thoughts as to how these may be addressed going forward. I also discuss several areas of future work, highlighting where I believe this work will have the greatest impact.

## 1.2 List of publications and author contributions

The contributions for all papers are the same. All papers are co-authored by the supervision team (Tianning Tang, Stephen Roberts, and Thomas Adcock) except for the first chapter which is only TT and TA. The supervision team provided guidance and steering on the technical work. All the conceptualization, coding, detailed model development, analysis and initial drafts were done independently by myself. A list of the manuscripts included in this work is as follows:

1. Tidal corrections from and for SWOT using a spatially coherent variational Bayesian harmonic analysis (Published: JGR Oceans) (13)
2. Observations of the seiche that shook the world (Accepted Nature Communications) (14)
3. Automating the tidal response method (Accepted: JGR Machine Learning and Computation) (17)
4. Response-based prediction of tidal currents (Under Review: JGR Oceans)
5. A hybrid model for Online short-term tidal energy forecasting (Published in Applied Ocean Research (20))

# 2 Literature Review

## 2.1 Tides

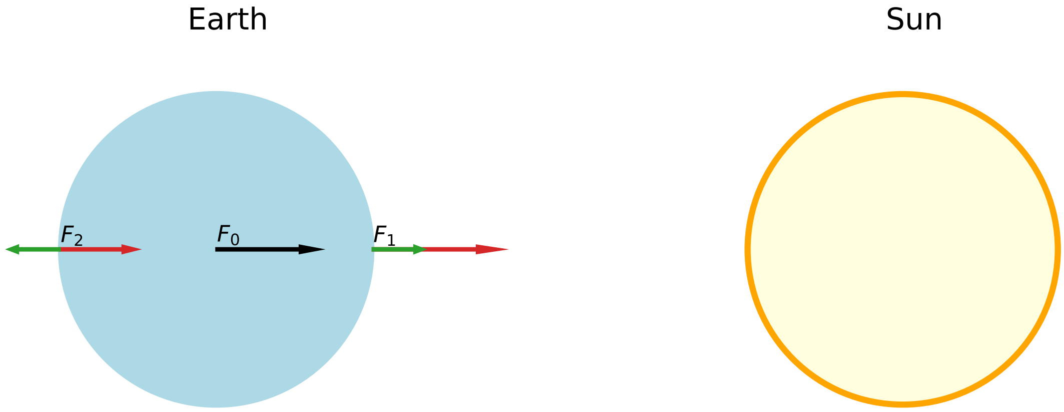

Figure 2.1: **Two body example of the gravitational forcing exerted by the Sun on the Earth.** $F_0$ is the force acting on the center of mass. The red arrows $F_1$ and $F_2$ are the forces acting on the opposite sides of the Earth. The green arrows represent the differential force relative to the center of mass $F_{diff} = F_0 - F_N$.

Tides are the periodic movements of oceans due to gravitational forcing from the combined Earth-Moon-Sun system. In this context, tidal heights and currents refer to the vertical and horizontal components of these movements respectively. A simple intuition of tides can be built directly from these forces. Let us start with a simple example with only two bodies, taken to be the Earth and either of the Moon or Sun. For simplicity we will treat the Earth as a frictionless water planet. As will be discussed, continents, friction, irregular and finite depths, turbulence, and other features will induce deviations of the actual tides from our toy model. For two bodies with mass $M_1$ and $M_2$ with distance between center of masses $R$, Newton's law of gravitation defines the force on the center of mass as

$$F_0 = G\frac{M_1 M_2}{R^2}, \tag{2.1}$$

where $G = 6.67410^{11} Nm^2/kg^2$. This force is denoted by the black arrow in Figure 2.1. Let us now consider two representative points $F_1$ and $F_2$ on opposite sides of the Earth. These points represent the closest and furthest points from the 2nd celestial body. The force on each point can once again be computed with $F_1 = G\frac{M_1 M_2}{(R-r)^2}$ and $F_2 = G\frac{M_1 M_2}{(R+r)^2}$ where $r$ is the radius of the Earth. The corresponding forces are shown by the red arrows in Figure 2.1. We can compare these two forces from the reference frame of the center of mass of the Earth by subtracting $F_{diff} = F_0 - F_N$. After some algebraic manipulation this yields the gravitational differential with

$$F_{diff} = -2GM_1 M_2 r/R^3 \tag{2.2}$$

which is shown by the green arrows in Figure 2.1. It can be seen that the gravitational differential produces two opposite forces on either side of the Earth with the point closest to the celestial body pointing towards its center of mass, and the point furthest pointing the opposite direction. It is this differential which drives global tides, and is the reason why a secondary bulge is seen opposite the gravitational force.

This is not the tidal bulge but rather the envelope of the total tidal potential, which defines the equilibrium tide. In our toy example, this is equivalent to the bulge which is the actual deformation of sea-level to these forces. In this context, if we consider the Sun as our celestial body, then any given location will experience two high tides and two low tides a day as the Earth rotates through each bulge which remain aligned with the Sun. In the full system comprising the Earth, Moon, and Sun, both a lunar and solar potential are induced creating two different bulges. The superposition of these two bulges gives rise to the Spring-Neap tidal cycle whereby maximum (Spring) tides occur when the Moon and Sun are in alignment, and minimum (Neap) tides occur when they are perpendicular to one another. Despite its much smaller mass, the moon actually exerts a greater control over Earths tides (approximately 2 times greater) due to its relatively close proximity as compared to the Sun.

While the tide-generating forces produce dual bulges on each side of the Earth, there are not necessarily two tides each day everywhere. The tilt of the Earth on its axis (approx 23.5 degrees) leads to asymmetry in the magnitude of the tidal bulge at various latitudes. Given this axis of rotation, a diurnal inequality in the magnitude of semi-diurnal high-tides will be observed at locations above and below the equator. In some regions this inequality is so great that only one high-tide is seen per day, creating so-called diurnal tides.

Further complexities arise as a consequence of the movements of the combined Earth-Moon-Sun system. In both Earth-Moon and Earth-Sun systems, the two objects orbit around a common center of mass, producing elliptic orbits. The variation in distance between objects during such orbits also impacts the magnitude of the tide-generating force as shown. These cycles occur over a solar year (approx 365.2596 days) and a lunar month (approx 27.554 days), respectively, giving rise to what are called long period tides.

The final major cycle of interest is created by the tilt of the Moon's orbit by approximately 5.1 degrees relative to the ecliptic plane (the plane of Earth's orbit around the Sun). As a result, the lunar declination — the angular distance of the Moon North or South of the Earth's equator — varies over time. This variation modulates the amplitude and latitude of the tidal bulges, producing a long-period oscillation known as the nodal cycle. The full cycle of this variation, has a period of approximately 18.6 years. Nodal modulations are observable in most long-period observations and thus play a key role when evaluating changes in sea-level over longer periods.

### 2.1.1 Describing Tides

The cyclic nature of the tide-generating forces gives rise to a natural procedure for their analysis and description: the decomposition of the tide into its component frequencies, known as *harmonic analysis*. The mathematical basis of harmonic analysis is described in Section 2.2.2. Here, I introduce the basic terminology used to describe tides.

In harmonic analysis, each frequency of the tide-generating potential produces a *harmonic constituent*. These are denoted by a combination of letters and numbers, such as $M_2$, $S_2$, and $O_1$. The letters indicate the astronomical source (e.g., $M$ for Moon, $S$ for Sun, $K$ for a lunar-solar interaction), while the subscript refers to the periodicity in cycles-per-day: '1' for diurnal (approximately 24-hour), '2' for semi-diurnal (approximately 12-hour), and so on.

Additional constituents arise due to nonlinear effects such as shallow water interactions and friction. *Compound constituents* result from the interaction of two or more astronomical constituents (e.g., $MS_4$ from $M_2{+}S_2$), while *overtides* are higher harmonics of primary constituents (e.g., $M_4$, the first overtide of $M_2$). These terms help capture tidal asymmetries and distortions in coastal regions.

Each constituent is represented by a sinusoid with a fixed frequency and fitted amplitude and phase. Their superposition reconstructs the tidal signal and allows for its prediction at a given location. This simple description informs the powerful harmonic analysis procedure

| **Species** | **Constituent** | **Period (h)** |
|---|---|---|
| Principal lunar semidiurnal | M2 | 12.4206012 |
| Principal solar semidiurnal | S2 | 12 |
| Larger lunar elliptic semidiurnal | N2 | 12.65834751 |
| Larger lunar evectional | $\nu$2 | 12.62600509 |
| Variational | $\mu$2 | 12.8717576 |
| Lunar elliptical semidiurnal second-order | 2N2 | 12.90537297 |
| Smaller lunar evectional | $\lambda$2 | 12.22177348 |
| Larger solar elliptic | T2 | 12.01644934 |
| Smaller solar elliptic | R2 | 11.98359564 |
| Shallow water semidiurnal | 2SM2 | 11.60695157 |
| Smaller lunar elliptic semidiurnal | L2 | 12.19162085 |
| Lunisolar semidiurnal | K2 | 11.96723606 |
| Lunisolar diurnal | K1 | 23.93447213 |
| Lunar diurnal | O1 | 25.81933871 |
| Lunar diurnal | OO1 | 22.30608083 |
| Solar diurnal | S1 | 24 |
| Smaller lunar elliptic diurnal | M1 | 24.84120241 |
| Smaller lunar elliptic diurnal | J1 | 23.09848146 |
| Larger lunar evectional diurnal | $\rho$ | 26.72305326 |
| Larger lunar elliptic diurnal | Q1 | 26.86835 |
| Larger elliptic diurnal | 2Q1 | 28.00621204 |
| Solar diurnal | P1 | 24.06588766 |
| Lunar monthly | Mm | 661.3111655<br>27.554631896 |
| Solar semiannual | Ssa | 4383.076325<br>182.628180208 |
| Solar annual | Sa | 8766.15265<br>365.256360417 |
| Lunisolar synodic fortnightly | MSf | 354.3670666<br>14.765294442 |
| Lunisolar fortnightly | Mf | 327.8599387<br>13.660830779 |
| Shallow water overtides of principal lunar | M4 | 6.210300601 |
| Shallow water overtides of principal lunar | M6 | 4.140200401 |
| Shallow water terdiurnal | MK3 | 8.177140247 |
| Shallow water overtides of principal solar | S4 | 6 |
| Shallow water quarter diurnal | MN4 | 6.269173724 |
| Shallow water overtides of principal solar | S6 | 4 |
| Lunar terdiurnal | M3 | 8.280400802 |
| Shallow water terdiurnal | 2MK3 | 8.38630265 |
| Shallow water eighth diurnal | M8 | 3.105150301 |
| Shallow water quarter diurnal | MS4 | 6.103339275 |

Table 2.1: Table of important tidal constituent, their origin, and period.

described in Section 2.2.2 and developed further in this thesis. As will be discussed, while this formulation has largely driven the language we use to describe tides, it is not the only lens that can be used to describe or predict them. This distinction will be explored extensively in the following section. A table containing commonly encountered tidal constituents discussed in this thesis is in Table 2.1

## 2.2 Tidal Analysis and Prediction

The origins of tidal analysis and prediction can be traced back centuries to early tide tables in the UK from the 1200's (10). Indeed, one could devote an entire thesis to the history of tides. While the timeline of these methodological developments is interesting, it is simpler to start any discussion of tides from the governing equations. As will be shown, this provides a natural point of departure to the classical methods of tidal analysis and prediction. This review is meant to serve as an introduction to the strengths, weaknesses, and assumptions of the classical methods of tidal analysis in order to motivate the methodological developments introduced in this thesis. Comprehensive presentations of the theory of each approach are given in the respective chapters where appropriate.

Due to the length of the tidal waves, tides behave as *shallow-water waves.* Their movement is thus governed by the shallow-water equations. Consider a Cartesian reference frame with $x$, $y$ and $z$ representing the East-West, North-South and vertical directions, respectively. Let the mean-sea-surface correspond to $z = 0$. Adopting a notation similar to (16) the dynamic equations and associated continuity equation are given by

$$\begin{cases} \dfrac{\partial \bar{u}}{\partial t} + \dfrac{\partial \bar{u}^2}{\partial x} + \dfrac{\partial \bar{u}\bar{v}}{\partial y} - f\bar{v} = -g\dfrac{\partial}{\partial x}(\zeta - \zeta_g' - \zeta_p') + \dfrac{1}{\rho h}\left(\tau_x^w - \tau_x^b\right), \\ \dfrac{\partial \bar{v}}{\partial t} + \dfrac{\partial \bar{u}\bar{v}}{\partial x} + \dfrac{\partial \bar{v}^2}{\partial y} + f\bar{u} = -g\dfrac{\partial}{\partial y}(\zeta - \zeta_g' - \zeta_p') + \dfrac{1}{\rho h}\left(\tau_y^w - \tau_y^b\right), \\ \qquad\qquad \dfrac{\partial \bar{u}}{\partial x} + \dfrac{\partial \bar{v}}{\partial y} = -\dfrac{1}{h}\dfrac{\partial \zeta}{\partial t}. \end{cases} \tag{2.3}$$

Here, the overbar indicates the depth-average of the given quantity and:

- $u, v$ = velocity components in the $x$ and $y$ directions (m/s),
- $h(x, y)$ = bottom topography such that depth $z = -h(x, y)$

- $\zeta(x, y, t)$ = surface elevation
- $\zeta'_g$ = Equilibrium tide
- $\zeta'_p$ = Inverse barometric pressure effect
- $\rho$ = density of seawater (kg/m$^3$),
- $g$ = gravitational acceleration (m/s$^2$),
- $\tau^w_x, \tau^w_y$ = wind stress components in the $x$ and $y$ directions (N/m$^2$),
- $\tau^b_x, \tau^b_y$ = bottom stress components in the $x$ and $y$ directions (N/m$^2$)
- $f$ is the Coriolis force

These basic equations are the foundation on which all methods of tidal analysis and prediction are built. Using this starting point, I here detail the basis of these methods and the assumptions they make.

### 2.2.1 Numerical Prediction

The most obvious method of analysis given Equation 2.3, is to directly solve these governing equations. Given the ill-posed nature of solving these partial differential equations directly, these solutions are instead approximated by discretizing a domain of interest such that simplified (and analytically tractable) versions of the governing equations can be applied point wise. Numerous variants of this approach exist; however, the leading approaches rely on finite element solutions (22). There are three key components which determine the accuracy of the given model:

First, and most important, is the accuracy of the model's input conditions. Here, I use *input conditions* as an encompassing term for all user defined information e.g. bathymetry, forcing, drag parametrization etc. The impact of each of these inputs is readily seen in Equation 2.3. The nature of these inputs depends on whether the given model is regional or global. For regional modeling, tidal forcing is often not explicitly included in the model, but rather imposed implicitly by fixing tidal elevations from linear tides on the oceanic boundary. This procedure is justifiable as the local tide generating forces in shelf-seas are dominated by forcing from the open ocean boundary. However, this does produce notable uncertainties if the reference tidal model is inaccurate, or only a subset of tidal constituents are used as forcing. In contrast, global models are forced by the gravitational potential and thus do not posses this source of uncertainty (23). The selection of bottom drag parameters (and indeed accurate bathymetry) has significant implications for the associated tidal currents and tidal propagation in estuarine regions as $\frac{1}{h}$ becomes greater. Conventional practice is to 'tune' drag parameters based on observations. This becomes problematic when no such observations exist and injects a degree of empiricism into these 'physical' models. Once defined, input conditions such as bathymetry and drag coefficients are conventionally regarded as static. While this assumption is generally sufficient, problems can arise due to shifting sandbanks from large storms.

Second, the model resolution dictates how well small-scale processes and topographic features are captured (24). Higher-resolution models can resolve finer details of coastal geometry, and more localized effects leading to more precise predictions. However, increased resolutions significantly increase the computational cost presenting severe bottlenecks when modeling large domains and in global models (25).

The third problem with numerical modeling actually lies in Equation 2.3. One of the most important contributions of satellite altimetry has been the discovery that significant percentages of tidal energy are dissipated in deep oceans due to internal tide conversion (26). This conversion refers to the transfer of barotropic tidal energy into internal (baroclinic) waves as tides interact with rough topography, leading to energy dissipation in the ocean interior. As written, Equation 2.3 neglects such processes. This can be overcome by either parameterizing the dissipation within the 2-D shallow-water equations (SWE), or running a full 3-D model of the stratified ocean. The former is less accurate, and the latter incurs extreme computational expense.

While numerical modeling is 'physical', in the sense that it is solving the dynamic equations of motion, inaccuracies in the associated domain or boundary conditions can yield a-physical solutions. Hence, we encounter the same problems of empirical modelers – that is, the quality of the numerical solution is only as good as the input data. A salient recent example of this was the laborious development of the new UK Surge NEMO model, which increased resolution from 7 km to 1.5 km (27). Surprisingly, the model's performance improvements are mixed, and in some cases, worse. While increased resolution helps the resolution of small-scale processes, if the features e.g. bathymetry which produce such features are incorrectly represented such increases can actually do more harm than good.

Due to limitations in grid resolution for global models, some regions globally will inevitably go unresolved. To combat this, standard practice is to perform simple interpolation within these regions. Due to the strong dependence on $h(x, y)$, in Equation 2.3, this introduces considerable uncertainty in complex coastal regions. The extent of the errors this produces has only recently become possible to quantify with new wide-swath altimetry missions (28).

#### 2.2.1.1 Assimilative Models

Given the aforementioned limitations of purely hydrodynamic solutions, assimilative approaches have been developed to use observations to constrain the model solutions. These approaches have been highly successful for global tidal modeling owing to their ability to 'over-ride' regions with suspect bathymetry (22; 29; 23). The basic idea of these approaches is to produce solutions which yield an optimal balance between the observed data and the adherence to the hydrodynamics. The addition of this constraint introduces even greater computational overheads, particularly when using time-domain solvers (22). As such, modern assimilative tide models employ 'Spectral Ensemble Optimal Interpolation' which is a variational approach based on the representer method (30). Here, tidal harmonics which have been computed from in-situ tide gauges and the decades of available altimetry data (as described at length in the following Section 2.2.2), are assimilated directly within the frequency domain. This presents a trade-off between predictive accuracy and the correctness of downstream scientific products such as energy diagnostics. Furthermore, the assimilation of vast swaths of available data brings into question the validation of these models. For example, the Finite Element Solution (FES) produced by CNES assimilates a majority of the global tide-gauges subsequently used for validation (23), and does not publicly list the tide gauges employed. These issues are explicitly addressed in Chapter 3 including quantifying the spatial errors within these models, and development of a metric which can quantify the undue biasing from validating on assimilated data. An extended description of FES is provided at the end of this chapter.

#### 2.2.1.2 Modern Usage and Challenges

Numerical and assimilative modeling is ubiquitous with regards to global and regional tide and storm surge predictions (23). This is particularly true for modeling tidal currents, for which observational data is sparse and expensive to obtain. The physical basis of these methods ensures solutions are consistent and explainable; both of which are seen as important for operational applications. As such, to the best of my knowledge, numerical/assimilative approaches are the only methods employed operationally for storm surge forecasting[1].

However, as noted above major challenges still remain. First, while these models are reliable due to their physical basis, they can be reliably inaccurate if the input conditions are inaccurate. This can be seen by the superiority of empirical tidal models in coastal regions (19; 23). Additionally, new wide-swath altimetry missions such as SWOT are producing exponentially more data than classical altimetry. As such, the usage of these data in assimilative models is presently intractable.

### 2.2.2 Harmonic Analysis

The harmonic method was introduced by Sir William Thompson (Lord Kelvin) in 1867 (31) and modernized by Doodson (11) and countless others over the next 150+ years (32; 33; 34). As an empirical approach, the harmonic method deals exclusively with the observables of Equation 2.3. Based on the observed sea-level $\zeta$ (or current measurements in the case of tidal currents), we seek to construct a model $\hat{\zeta}(t)$ with parameters $\theta$ such that

$$\arg\min_{\theta} \frac{1}{T} \sum_{t=1}^{T} \left(\zeta(t) - \zeta'(t)\right)^2 . \tag{2.4}$$

To do this we must assume an appropriate form of the model $\hat{\zeta}$. Due to the periodic nature of tides, harmonic analysis (HA) assumes the tidal signal can be decomposed into a series of sinusoidal waves at frequencies given by the forcing which produces them. Mathematically, the

---

[1] Notably, the method described in Chapter 5 is operationally deployed at the Thames Barrier in London and in the Dutch operational HMC system and has greatly improved forecast accuracy over these models.

tidal signal $\hat{\zeta}$ can be related to the observed sea-surface height $\zeta$ by

$$\zeta(t) = \hat{\zeta}(t) + \epsilon(t) = \sum_{k=1}^{N} C_k \cos(\omega_k t + \phi_k) + \epsilon(t), \tag{2.5}$$

where $C_k$, $\omega_k$, and $\phi_k$ denote the amplitude, angular frequency, and phase,of the $k$-th tidal constituent, respectively and $\epsilon(t)$ is the non-tidal residual. The discrete frequencies $\omega_k$ can be derived directly from the tide-generating potential as in (11). The choice of which constituents to include in a given analysis is of significant importance and is discussed at length in Chapter 3.

It is worth evaluating how the above two Equations fit into Equation 2.3. First, we assume the sea-level can be decomposed into tidal and non-tidal signals, with the tidal signal being that which is produced by the gravitational forcing alone. This leads to a simplified momentum balance of the form:

$$\begin{aligned} \frac{\partial \bar{u}}{\partial t} &\approx -g\frac{\partial}{\partial x}\left(\zeta - \zeta_g'\right), \\ \frac{\partial \bar{v}}{\partial t} &\approx -g\frac{\partial}{\partial y}\left(\zeta - \zeta_g'\right). \end{aligned} \tag{2.6}$$

The residual terms relating the atmospheric pressure, advection, shallow-water effects (depth dependence of bed-friction and wind), and bottom friction terms are thus considered to be part of $\epsilon(t)$. Notably, wind stress, atmospheric pressure, bottom friction, and nonlinear advection all appear in the shallow-water equations and are dynamically coupled with the tides. The last three terms are nonlinear and can directly modulate the tidal signal itself. If the tidal signal existed in isolation, e.g. no nonlinear modulation existed, the assumption that the signal can be represented by a harmonic expansion would hold. The problems with this formulation come into play when we try and estimate the associated tidal constituent's amplitudes and phases. For modern harmonic analysis, this is done in the time-domain using least-squares based estimation. To do this, the harmonic analysis is cast as a linear model $\hat{\zeta}(t) = w^{\mathsf{T}}x + \epsilon(t)$, where $w$ is a weight vector of derived quadrature amplitudes and $x$ is a design matrix of the form

$$x = [\sin\omega_1 t, \cos w_1 t, \ldots, \sin\omega_k t, \cos w_k t]. \tag{2.7}$$

Casting the model in this way and solving with least-squares necessitates that we follow its rules. Specifically, least-squares imposes assumptions that the residuals $\epsilon(t)$ are independent, identically distributed, and normally distributed with constant variance (homoscedasticity). If these assumptions are violated—such as when the residuals exhibit autocorrelation or heteroscedasticity due to unmodeled nonlinear dynamics—the resulting estimates of the tidal constituents' amplitudes and phases may be biased or inefficient, and the standard error estimates become unreliable.

#### 2.2.2.1 Modern Usage and Challenges

Harmonic analysis is the cornerstone of modern tidal analysis. The approach is utilized not just for routine analysis and prediction tasks, but has also defined the language used to discuss and study tides. Its simple and automated usage provided jointly by least-squares estimation techniques and automated constituent selection criterion have made it accessible to non-expert users.

While in a majority of cases the assumptions of harmonic analysis are tolerable, several major challenges remain, predominantly in coastal and estuarine environments. In estuarine environments, the tidal signal can become distorted due to nonlinear interactions with the bathymetry, friction, and hydrodynamic effects. This distortion leads to asymmetries in the tidal signal, such as saw-tooth profiles, which are not well represented by a pure harmonic approach (35). Further issues occur in areas of wetting and drying where "discontinuities" in the tidal signal yield significantly underestimated harmonics (28). Per Equation 2.3, it is no surprise that fluvial and meteorological effects can also severely bias harmonic analysis. With regards to the former, the nonstationarity in tidal rivers has necessitated the development of new methods as described in Section 2.2.4. Meteorological effects on the other hand are not easily accounted for in a harmonic framework and most work has necessarily gone into developing estimators which reduce their impact on the derived constituents (33; 34). However, even by modifying the assumptions about the noise $\epsilon(t)$, significant biasing of the harmonic constituents can still occur (36).

The second major limitation of harmonic analysis is the requirements it imposes on time-series duration and sampling. Whilst not a physical-law, the Rayleigh criterion rightly asserts that the resolution of two tidal frequencies $f_1$ and $f_2$ is related to the length of the tidal record

$T$ by $\frac{1}{T} \le |f_1 - f_2|$. In reality, this is also a function of the noise in the signal (1). Further complications are introduced with irregular temporal sampling and 'aliasing'. Typically, discussions of aliasing are centered around the Nyquist-Shannon sampling theorem which states a frequency $f$ can only be resolved unambiguously if the sampling interval $\Delta t$ satisfies $f < \frac{1}{2\Delta t}$ (37). This is not necessarily correct as the tidal signal has a sparse Fourier basis, and thus accurate constituent estimations can be obtained below the Nyquist frequency (38). However, as will be shown in Chapter 5, a third (and previously unconsidered) component is the estimator employed for analysis. This can alleviate some of the above limitations.

### 2.2.3 Response Method

Whereas HA assumes the oceanic response to be strictly periodic, the response method assumes the response is either linear, or weakly nonlinearly related to the input forcing (5; 16). Under this assumption, our model learns a set of weights $w$ which relate the input forcing $V(t)$ at various lead-times $\tau$ to the observed sea-level $\zeta$ through the convolution:

$$\hat{\zeta}(t) = \sum_s w(s) V(t - \tau_s) \tag{2.8}$$

with $\hat{\zeta}(t)$ denoting the predicted sea-level. Munk and Cartwright were the first to demonstrate that solutions of this form hold across a broad range of tidal regimes (5) and under non-tidal forcing (16). In this formulation, the assumptions made about Equation 2.3 are determined by both the input forcing, and how these inputs are allowed to interact. With regards to tides, the assumptions of the model are identical to those in HA (Equation 2.6), that is with the exception of the functional form of the solution. Countless works have shown response approaches to yield slightly more accurate analyses (5; 39; 40; 41). Furthermore, by avoiding separating the tidal spectrum into discrete frequencies, one can generally get away with using fewer parameters and shorter reference series as will be shown in Chapter 6.

Applications which include additional forcing inputs (beyond tidal forcing) have been limited outside of Cartwright's original work (16). While adding in these effects can reduce the simplifying assumptions about Equation 2.3 – and even yield superior estimates of the true tidal signal in the process – it also carries with it the obligation to correctly define how such forcing interacts with and modulates the tidal signal. This introduces a model selection problem into the estimation problem. Naturally, this grows more difficult as the interactions become more nonlinear and challenging if the nature of these interactions is poorly understood.

For example, suppose a non-stationary non-tidal force $G(t)$ interacts with the tidal signal $\zeta_T$ such that the total sea-level $\zeta(t) = F(\zeta_T, G; t)$. In a least-squares regression framework, omitting $G$ introduces bias in the estimated tidal coefficients due to the modulation of $\zeta$ by $G$. Specifically, if the interaction is nonlinear or temporally structured, the least-squares estimator will misattribute some of the non-tidal variability to the tidal signal, leading to systematic distortions in the estimated amplitudes and phases.

Incorporating $G$ into the regression model offers a means to mitigate this bias, contingent on correct specification of $F$. If $F$ is perfectly specified, the model can achieve zero error, meaning perfect separation of tidal and non-tidal components. However, this also introduces the possibility of exacerbating errors: an incorrect specification of $F$ can lead to overfitting, where spurious correlations drive the regression coefficients away from their true values. This presents a classic model selection tradeoff—while additional forcing terms can improve fidelity, they must be chosen with care to avoid introducing larger biases than those present in the simpler, tide-only model. Because of this, the incorporation of non-tidal forcing is perhaps the least-used aspect of the response method; however, as Chapter 5 will show, it has tremendous potential if it can be used.

#### 2.2.3.1 Modern Usage and Challenges

While many ideas within the original response method paper have become ubiquitous (smoothness of oceanic response, admittance functions, tidal cusps, etc.), the response method itself has been employed in comparatively few contexts. The primary application(s) is in satellite altimetry where the frequency resolution benefits are particularly useful for handling the sparse and strongly aliased sampling (39; 40). Notably, while these models demonstrate significant advantages over harmonic methods for linear tides, nonlinear tides are not estimated.

Indeed, the biggest challenge when employing the response method is the estimation of nonlinear responses. Neither the lumped or sequential procedures proposed in (42) can be automatically applied. This barrier has largely contributed to the rise of the harmonic method as the basis of tidal analysis and prediction. Chapter 5 and 6 solve this issue enabling the exploitation of the superior characteristics of the response approach.

Another shortcoming of the response method is the difficulty of interpretting the learned response weights. Unlike harmonic constituents which correspond to waves at tidal frequencies, response weights define a transfer function—the so-called admittance—from the forcing to the response. This admittance has considerable utility for inferring smaller constituents from major ones, if the condition of smoothness holds. However, in coastal regions with large resonances, admittances are not necessarily smooth which can undermine standard response analysis (43). This challenges are further compounded by the fact that the weights actually lose physical meaning when considering non-tidal contributions (16). These challenges are similarly addressed in Chapter 5.

### 2.2.4 Other Nonstationary Approaches

There are many processes which are tidally driven, but may not be stationary or adhere to strict tidal periodicities. Tidal rivers and tides under meteorological forcing (storm surge) are simple examples of this. Due to the fact that these processes violate the core assumption of harmonic analysis, specialized approaches have been developed. To my knowledge, the only approach which has routinely been employed for the prediction of tidal rivers is NS_Tide (15). This approach starts from the familiar starting point of HA, but replaces the conventionally static parameters $b_0, C_k, \phi_k$, with functions of the river outflow $Q(t)$ and diurnal tidal range $R(t)$.

A limitation of NS_Tide is the requirement of the user to define physically realistic basis functions. While this has been developed for riverine tides, laborious analytical development is required if one wishes to apply the approach to different physical processes. Alternatively, empirical analysis can be conducted using wavelet based approaches (44). These approaches yield a linear analysis procedure which can cleanly separate tidal and non-tidal variance. Unfortunately, while wavelet approaches yield remarkable analysis capabilities, prediction is generally unsatisfactory as there is no explicit connection between external forcing and the non-stationary behavior learned (45). Due to these limitations, I am not aware of any context in which wavelets are employed routinely for prediction. Hence, I do not consider this approach further.

## 2.3 Data

Any discussion of empirical methods of tidal analysis must be done in conjunction with the data employed for analysis. As described above, the duration and sampling characteristics of the data and presence of noise impacts the efficacy of the tidal analysis. A brief overview of the three types of observational modalities considered in this thesis is given below.

### 2.3.1 Satellite Altimetry

Satellite altimetry provides a globally distributed dataset of sea level observations, offering broad spatial coverage that is particularly valuable in regions where in-situ measurements are sparse or logistically infeasible (46). Unlike tide gauges, which provide high-temporal-resolution records at fixed locations, altimetry data are characterized by intermittent sampling dictated by satellite repeat cycles, typically ranging from 10 to 35 days. This temporal sparsity limits the resolution of high-frequency tidal constituents and can introduce aliasing effects, which must be carefully addressed in tidal harmonic analysis (47).

Additionally, historical altimetry measurements degrade in the coastal zone due to land contamination within the instrument footprint, increased noise from variable sea states, and uncertainties in corrections for atmospheric and geophysical effects. While new missions such as SWOT have partially addressed these errors (48), a range of errors still exist and are discussed and addressed in Chapter 3.

### 2.3.2 In-situ Water-level observations

In-situ water-level observations have been obtained for centuries using a range of measurement techniques (49). These observations are typically high-frequency (hourly or less), and have relatively low error (as compared to altimetry). While the instruments employed for this task have improved drastically, similar considerations remain when using these data due to the range of measurement errors present. A brief list of considerations is as follows:

1. **Time-shifts** in which the reference time is shifted systematically, create problems for tidal harmonic analysis. This leads to inaccuracies in derived phases of constituents and is not readily detectable through statistical outlier identification procedures.
2. **Biofouling and dropout** occur when biological growth obstructs sensors or when gauges

temporarily fail, leading to missing or unreliable data. Biofouling can dampen high-frequency variability, distorting tidal amplitude estimates, while dropout results in gaps that degrade harmonic analysis, especially for long-period constituents. These issues complicate comparisons with altimetry and require careful quality control measures to mitigate their impact.

3. **Incorrect geolocation** can manifest itself in two forms as gauges are referenced to both vertical datums and fixed geolocations. Ground motion can shift gauges relative to the vertical datum which can create problems for sea-level trends (50). SWOT data has also revealed inaccuracies between reported and actual gauge geolocations. This creates problems when assessing the performance of tidal models and altimetry.

#### 2.3.2.1 In-situ current observations

In general, in-situ oceanic current measurements are more expensive to obtain than water-levels, and are consequently much shorter (10). Naturally, this is problematic when performing harmonic analysis as it limits the number of constituents which can be accurately estimated. Handling these shortened reference series is discussed explicitly in Chapter 6. Additionally, wave-action and biofouling can lead to systematic noise artifacts which create major problems for any type of analysis. Historically, current meters were placed on surface buoys, and were thus highly sensitive to surface wave contamination. Bottom mounted ADCPs, as employed in Chapters 6 and 7, mitigate these effects partially through spatial and temporal averaging.

It goes without saying that these devices can exhibit the same systematic errors as water-level counterparts. However, due to the considerable spatial variability between currents, and the directional nature of currents, problems can arise if a measurement device is moved or rotated between deployments.

### 2.3.3 Key Data and Benchmarks Considered

Now that the fundamentals of the methods and data employed for tidal analysis have been covered, I will briefly highlight some of the data sources and benchmarks which will be used throughout this thesis.

**GESLA 3** The Global Extreme Sea Level Analysis (GESLA) Version 3 dataset provides quality-controlled, hourly sea level observations from tide gauges around the world (51). It is a benchmark for validating tidal and sea level variability globally. The data comprise 5199 individual tide gauge records of hourly (or higher) temporal resolution. A limitation of these data is that they are concentrated in the Global North which can lead to biased estimates of global variability. The GESLA 3 data can be found here[2].

**TICON 3** The TICON (Tidal Constituents) Version 3 database offers a globally consistent set of harmonic tidal constituents derived directly from the GESLA 3 tide gauge records (52). Due to a pre-filtering process which removes gauges missing more than 70% of measurements, the TICON 3 data provides tidal estimates for 3,471 of the GESLA 3 gauges. It is used to benchmark tidal predictions and assess model skill in reproducing observed tidal harmonics. It is worth noting that these harmonic constituents are estimated using a basic least-squares harmonic analysis. Care should thus be taken in considering these to be the ground truth. Furthermore, these data inheret the same geographic biasing of GESLA 3. The TICON 3 data can be found here[3].

**Surface Water Ocean Topography Mission (SWOT)** The NASA Surface Water and Ocean Topography (SWOT) mission is the first altimetry satellite to provide two-dimensional maps of sea surface height at unprecedented spatial resolution (10s of meters) across a wide swath (120km) (53; 48). During its 102-day calibration and validation phase, SWOT operated in a fast-repeat orbit, offering near-daily temporal sampling—unlike any previous altimetry mission—which is favorable for studying tidal and subtidal signals. As a result, SWOT enables us to evaluate tides in a fundamentally new way, particularly in coastal, shelf, and estuarine environments where existing models typically perform poorly. It also allows for further development of tidal analyses to take advantage of these unique characteristics. This thesis will consider data from both the calibration and validation phase, as well as the subsequent science phase. Further technical details of SWOT's orbit, sampling characteristics, and data products are provided in the relevant sections of this thesis. The data can be accessed through aviso[4].

---

[2] `https://gesla787883612.wordpress.com/downloads/`

[3] `https://doi.pangaea.de/10.1594/PANGAEA.896587`

[4] `https://www.aviso.altimetry.fr/en/missions/current-missions/swot.html`

**FES (Tide Model)** The Finite Element Solution (FES) tide model is the state-of-the-art global barotropic tide model based on data assimilation of satellite altimetry and in situ observations produced by CNES (22; 54). Several iterations of this model have been released over the years with the newest iteration, FES2022, being released after most of the in this thesis looking at global tide modeling was conducted. The difference between FES2022 and its predecessor FES2014b is the resolution of the bathymetric mesh and the number of gauges and altimetry points assimilated into the model. While other global tide models exist, FES tidal corrections are included as the standard correction for all altimetric data considered in this manuscript making it appropriate for benchmarking. The up-to-date version of FES can be accessed here[5].

---

[5] `https://www.aviso.altimetry.fr/en/data/products/auxiliary-products/global-tide-fes.html`

# 3 Tidal corrections from and for SWOT using a spatially coherent variational Bayesian harmonic analysis

# Abstract

The accuracy of global tidal models degrades significantly in coastal and estuarine regions. These models are important for correcting measurements from satellite altimetry and are used in numerous scientific and engineering applications. The new Surface Water Ocean Topography (SWOT) mission is providing measurements at unprecedented horizontal resolution in these regions. These data present both the opportunity and the necessity to quantify and correct the spatial variability in the model inaccuracies specific to these regions. We develop a variational Bayesian framework for tidal harmonic analysis which can be applied to SWOT, and is especially useful for exploting the data from the Cal/Val phase. The approach demonstrates superior robustness to different types of noise contamination in comparison to conventional least-squares approaches while providing full uncertainty estimation. By imposing a spatially coherent inductive bias on the model, we achieve superior harmonic constituent inference from temporally-sparse but spatially-dense data. Bayesian uncertainty estimation gives rise to statistical methods for outlier removal and constituent selection. Using our approach, we estimate a lower bound for the residual tidal variability for two SWOT Cal/Val passes (003 and 016) around the European Shelf to be 7% on average. We also show similar estimates cannot be produced using standard least-squares approaches. Tide gauge validation in the same region confirms the superiority of our empirical approach in coastal environments. Empirical corrections for the SWOT data products are provided alongside an open-source Python package, VTide.

## 3.1 Introduction

The NASA Surface Water Ocean Topography mission (SWOT), launched in December 2022, is currently producing the highest spatial resolution satellite altimetric measurements of global oceans to date (48). SWOT has the potential to greatly improve the field of coastal altimetry by providing high spatial and temporal resolution views into estuaries and coastal regions with measurements extending to within 50 m of coastlines. Whereas barotropic tides in deep oceans are essentially linear, interactions with coastal bathymetry, river outflow, and shallow water effects produce considerable nonlinearity. Due to data accuracy issues in historical altimetry and the increased nonlinearity present, the accuracy of global tide models tends to degrade in these regions with shelf and coastal errors in the derived M2 constituent ranging from 6-10% and 7-24% respectively (23). Although, improved tidal analysis and prediction is, of course, useful for studying tides, arguably more important is the accurate removal of tidal variability from altimetric observations. This is a consequence of the fact that the tidal signal often dominates other mesoscale and sub-mesoscale processes. These processes are of high interest to SWOT researchers. Developing improved methods of tidal analysis and prediction is therefore critical to both the success of the broader SWOT mission and in utilizing the decades of satellite altimetry data which already exists.

A unique feature of the SWOT mission was the 102 day calibration and validation phase. Characterized by a fast-sampling orbit with near 1-day repeats, this provided the highest temporal resolution of ocean altimetry missions to date, capturing 8 full cycles of the aliased M2 signal. In addition to the improvements in temporal resolution, SWOT data provide sea surface height measurements at unparalleled horizontal resolution using two Ka-Band interferometers (KaRIn). Combined, these swaths cover nearly 120 kms and provide SSH data products at a resolution of (250 m × 250 m) globally between $\pm 77.6^\circ$ latitude. This improved sampling can shed light on many short (temporal and spatial) scale processes such as internal tides, storm surge, and the coupling of river discharge and tides in estuarine regions (55). While this work does not make use of the 50 m × 50 m pixel cloud hydrology product, it should be noted that the methods developed herein can also be used on these data.

This work presents a variational Bayesian framework for spatially coherent tidal corrections which we apply to early stage ($\leq 1$ year) of SWOT data. Our approach can be applied to any location, without prior knowledge of bathymetry or gauge constraints. Furthermore, our fully Bayesian framework provides accurate uncertainty estimation, allowing researchers to evaluate how much a given result should be trusted. This uncertainty captures both the intrinsic uncertainty from the data due to noise (aleatoric), as well as the uncertainty in the derived parameters due to insufficient information (epistemic). Quantifying epistemic uncertainty is particularly important for sparse reference series where insufficient data is a primary source of uncertainty. We develop several methods for utilizing the uncertainty information to remove outliers and determine whether tidal constituents should be included in a given analysis. This approach makes the incorporation of prior information simple while still generating predictions from the SWOT data. These features enable harmonic constituent "super resolution" (estimating constituents from reference series shorter/more sparsely sampled than required by Rayleigh Criterion/Nyquist. See Section 3.2.6). The proposed framework can function both as a tool for direct regional tidal correction and as a means of assessing the tidal errors present in current models. The latter is critical to address the spatial limitations of shallow-water and coastal tide gauge datasets used for validating global tide models (56).

### 3.1.1 Empirical Correction Challenges and Philosophy

The unique spatiotemporal characteristics of the SWOT data and the challenges encountered when conducting tidal analysis on sparse data warrant the development of a specialized methodology. Here we look to briefly highlight some of these characteristics in order to motivate our methodological developments. The temporal sparsity intrinsic to satellite altimetry will always present a challenge for empirical correction of tidal phenomena due to aliasing of the harmonics. While SWOT's Cal/Val orbit produces favorable alias periods for semi-diurnal constituents, with its near 1-day repeats, this phase lasted only 102 days. Thus, at best we are left with just over a hundred usable measurements per location. In practice, however, we encounter a second difficulty when performing tidal analysis. While SWOT has extremely low measurement noise, early SWOT data has revealed several other sources of error (intertidal regions, roll errors, etc.), including the unavoidable "contamination" from non-tidal processes, especially storm surge (57; 58). Even tidal processes can give rise to errors as a consequence of aliasing at certain frequencies (particularly diurnal constituents for SWOT) into the mean. These contributions are amplified by the limited amount of data. The resultant biasing of the derived constituents can lead to significant errors. Hence, the problem of deriving tidal harmonics from early-stage ($< 1$ year) SWOT data is composed of two separate problems; true "aliasing" intrinsic to applying harmonic analysis to short and sparsely sampled reference series, and errors abounding from the interaction of tidal and non-tidal noise with the chosen estimator. While little can be done for the former, this paper demonstrates much can be done to address the latter.

The SWOT data has already been corrected with state-of-the-art tidal corrections from the assimilative finite element solution FES2014b (22; 59). Our empirical correction philosophy is to treat this as a baseline, and only make corrections if our model's confidence in the derived parameters exceeds a given probabilistic threshold. Thus, if our model's confidence in the correction at a given location sits beneath this threshold, we can simply default to the standard geophysical corrections. Further, we also consider a soft fusion of geophysical corrections with empirical model based values. Here, a soft fusion refers to combining these data probabilistically based on the model's confidence. Central to this philosophy is the ability to estimate when our model has gone wrong, which we accomplish by casting the problem within a Bayesian framework.

The objectives of this work are twofold:

1. Provide a framework to produce accurate uncertainty estimates for present tidal corrections.
2. Improve coastal and estuarine tidal estimates where possible.

The remainder of the chapter is structured as follows. Section 3.2 provides a detailed overview of our approach. Next, we extensively test and validate the proposed method in Section 3.3 and compare it to conventional methodologies using synthetic data. Section 4.4.1 then evaluates the proposed method on the L2 2 km SWOT data product using a combination of in-situ measurements from gauges and comparisons of sea-level anomaly (SLA) variance reduction. Here, the sea-level anomaly is the residual sea-level variability after all corrections (including tidal corrections) have been applied. Estimates of the residual tidal errors are computed and presented with confidence intervals to inform researchers where current corrections are deficient. Finally, Section 6.5 discusses the limitations of the proposed method alongside avenues of future work and Section 3.6 states conclusions.

## 3.2 Methods

### 3.2.1 Harmonic Analysis

Harmonic analysis is the cornerstone of modern tidal analysis and prediction. The largely coherent response of tides to the periodic astronomical forcing which gives rise to them makes this procedure highly effective and has enabled accurate tidal predictions for more than a century (10). Modern harmonic analysis is carried out in the time-domain using least-squares estimation (60), and in certain circumstances other weighted variants (33). To achieve this, the harmonic analysis is formalized as a general linear model such that $y_i = w^{\mathsf{T}} x_i + \epsilon_i$. Here, $y_i$ is the observed sea-level at time-step $i$, $x_i$ is the $i^{th}$ row of an $M \times N$ matrix of basis functions where $M$ is the number of measurements and $N = 2n + 2$ with $n$ equal to the number of constituents, $w$ is a set of inferred weights, and $\epsilon_i$ is the non-tidal residual. It is important to note that the assumptions made about the nature of $\epsilon$ have significant implications for the accuracy of the regression model used, which will be discussed later. The harmonic basis functions are defined within the design matrix $X$, with each column having the form

$$X = [1, t_i/t_{max}, \sin\omega_0 t_i, \cos\omega_0 t_i, \dots, \sin\omega_k t_i, \cos\omega_k t_i]^{\mathsf{T}} \tag{3.1}$$

where 1 is the mean-sea level, $t_i/t_{max}$ is the normalized linear sea-level trend, and the remaining entries corresponding to the quadrature basis functions for the $k^{th}$ constituent. A nuanced discussion on how constituents are selected is provided in Section 3.2.6.1. Mathematical details of nodal corrections are omitted from the paper for simplicity but are implemented in the standard way (61). It is generally preferable to solve for the quadrature amplitudes $A$ and $B$, though comparisons of tidal constituents are given in terms of the amplitude $C_k = \sqrt{A_k^2 + B_k^2}$ and phase $\phi_k = \arctan A_k/B_k$. We remark that $A_k$ and $B_k$ are elements of the weight vector $w$ corresponding to the $k^{th}$ set of quadrature basis functions. Note, the weights $w$ correspond to the parameters of the model, and are different from the weights used in weighted least-squares. It is these amplitudes and phases which serve as the basis of how we study and describe tides globally.

#### 3.2.1.1 Least Squares Parameter Estimation

Harmonic analysis constituent estimation is conventionally carried out using ordinary least squares (OLS). OLS estimation obtains parameter estimates which provide the most accurate description of the data in a least-squares sense. This imposes a strong prior assumption about the information content of the underlying data. In particular, conventional OLS assumes the residual term $\epsilon$ to be white Gaussian noise (WGN), with constant variance. In the case where this condition holds, the OLS parameter estimates are equivalent to those obtained via maximum-likelihood estimation (62) with weights $w$ given by:

$$w = (X^{\mathsf{T}}X)^{-1}X^{\mathsf{T}}Y, \tag{3.2}$$

where $Y = [y_0, y_1, \ldots, y_N]^T$ is a set of observations, and $X$ is the design matrix above. Problems arise in parameter estimation when the noise content of the data deviates from these assumptions. When the assumptions of least-squares are violated, harmonic constituent estimates which yield the most accurate estimation of the data are not necessarily equal to the true harmonics of the system. Countless oceanic, hydrological, and meteorological processes give rise to non-stationary and temporally correlated noise, which violate these assumptions (42; 63). These errors are propagated through the derived constituents and can yield significant biasing, such as the radiational double-counting problem (36). We note that while these issues have been partially addressed using weighted least-squares and other robust variants (64; 33), these approaches encounter different challenges due to the extreme sparsity of the early Cal/Val SWOT data and the need to select appropriate regularization parameters. Finally, perhaps the most significant limitation of conventional harmonic analysis procedures with OLS and one of high salience to satellite altimetry is the poor estimation of model uncertainty (19).

#### 3.2.1.2 Least-Squares Uncertainty Estimation

It is worth remarking under what conditions OLS uncertainty estimation is accurate. A comprehensive discussion and validation of least-squares uncertainty estimation techniques for tidal estimation is provided in (63), including IRLS estimators. Basic OLS confidence interval (CI) estimation is performed through comparison of the variance-covariance matrix of the regression such that

$$\mathrm{Var}(w) = \sigma^2 (X^\mathsf{T} X)^{-1}, \tag{3.3}$$

where $\sigma^2 = \sigma^2 I$ is the covariance matrix of errors given by

$$\sigma^2 = \frac{\epsilon\epsilon^\mathsf{T}}{M - N - 1}, \tag{3.4}$$

where M is the number of observations, and $N = 2n + 2$ is the number of estimated parameters with $n$ being the number of constituents here. These estimates are accurate when the following assumptions are satisfied

1. *Linearity* The relationship between dependent and independent variables is linear and yields a zero-mean residual.
2. *Independent Variables* There is no correlation between the estimated parameters and the errors.
3. *Homoscedasticity* Errors have constant variance
4. *No temporal autocorrelation* Errors are uncorrelated with each other.
5. *Gaussian Residuals* Errors are normally distributed.

As will be discussed in Section 3.2.6, the limited Cal/Val data necessitates that only a subset of constituents be used in a given analysis. Thus, constituents excluded from an analysis are now treated as noise by the model potentially violating assumptions 2. As is seen in Equations 3.3 and 3.4, the OLS standard error does not explicitly consider the dependence of $\epsilon$ on $X$ and can therefore lead to incorrect estimates. The uncertainty estimates computed in Figure 3.8 utilize this standard approach.

Several additional methods have been devised for confidence interval generation (32; 33; 34) and generally yield superior estimation in the presence of colored noise (63). However, due to the focus of this work on application to extremely sparse SWOT Cal/Val data these are not applicable. For example, testing the Monte Carlo (MC) confidence intervals (CI) in the UTide package yields singular-value decomposition (SVD) convergence errors. Similarly, the Linear CI estimation produces intervals containing NaNs due to the matrices used for inversion becoming rank-deficient (ill-conditioned) as the number of samples is reduced. In standard practice, these conditions are rarely encountered as the numbers of parameters is typically much smaller than the number of measurements.

Bootstrap methods, such as Moving Block Bootstraps and Semiparametric Bootstraps (only applicable to uniform sampling), generally yield superior uncertainty estimation to their analytical counterparts as they can better estimate the autocorrelation of the residuals (63). However, these approaches rely on separating residuals into many *sufficiently long* sub-series (blocks) of contiguous measurements. For the SWOT Cal/Val data, the short and sparse records are insufficient for this procedure and thus bootstrap methods are not considered further.

### 3.2.2 Bayesian Analysis

To combat the limitations discussed above and to achieve the objectives outlined in Section 3.1.1, we develop a fully Bayesian variant of harmonic analysis referred to herein as VBayes. The fundamental distinction between Bayesian and frequentist methodologies is the assumption that our parameters of interest are best represented by associated probability distributions, rather than as point values. This has the advantage that the uncertainty associated with the parameters in our model is explicitly represented. Numerous studies have shown that tides can be modulated by various forms of external forcing leading to smearing of the tidal spectrum (5) and variability in tidal constituents (15; 65). As such, we argue modelling tidal constituents as distributions rather than discrete lines is (i) more physically accurate and (ii) provides a more principled approach to handling this uncertainty. Under this assumption, we once again define the harmonic analysis problem as a general linear model $y_i = w^{\mathsf{T}} x_i + \epsilon_i$.

Using Bayes' theorem, we begin with a set of prior beliefs $p(\theta)$ about the distribution of all parameters (and hyper-parameters), which we denote $\theta$. It is worth at this point noting that a *hyper-parameter* is itself a parameter, but one that governs the probability distribution of a *parameter*. For example, the mean and variance of a Gaussian are the hyper-parameters which define the distribution over another variable. By extension, the term hyper-hyper-parameter is then used to describe a variable that controls the distribution over a hyper-parameter. Bayesian inference rarely moves beyond this level of complication. The task of Bayesian inference is to update prior beliefs over the distribution of $\theta$ to a posterior distribution, given a set of observations $Y = [y_0, y_1, \ldots, y_N]^T$ and the design matrix $X$. This posterior distribution, which describes our final beliefs over $\theta$, is given by:

$$p(\theta|Y) = \frac{p(\theta)p(Y|X,\theta)}{p(Y)}, \tag{3.5}$$

where $\theta$ represents the set of parameters and hyper-parameters of the model, $p(\theta)$ is the prior, $p(Y|X,\theta)$ the (data) likelihood, and $p(Y)$ the marginal likelihood, acting as a normalising term in the inference (as it does not depend upon $\theta$) (66).

**Likelihood:** The functional form of the (data) likelihood, $p(Y|X,\theta)$ in Equation 3.5, is consistent with the assumption of a least-squares error, thus is of Gaussian form. The squared residual, between observed $y_i$ and the model prediction, is weighted by hyper-parameter $\beta$, which represents a *precision*, or inverse (co)variance. This leads to a Gaussian likelihood term of the form (and noting that the likelihood only depends upon $w, \beta$):

$$p(Y|X,\theta) = p(Y|X,w,\beta) = \left(\frac{\beta}{2\pi}\right)^{N/2} \exp\left\{-\frac{\beta}{2}E_Y(w)\right\}, \tag{3.6}$$

where we denote $E_Y$ as the error functional, defined as:

$$E_Y = \sum_{i=0}^{N}(y_i - w^\intercal x_i)^2. \tag{3.7}$$

As described above, it is well documented that the tidal spectrum is super-imposed onto a red-noise background. To reduce the influence of this contamination, we introduce an automatic outlier determination procedure in Section 3.2.5 which exploits the assumptions of our noise prior. Testing using alternative noise priors, such as the red noise covariance defined in (67), while useful on high-frequency reference series, performed worse on SWOT Cal/Val data (0.99349 day sampling period). As such, this was not employed for this work.

Improving the handling of temporally correlated noise is an important area of future work, particularly in estuarine environments. In this vein, the VBayes framework can be modified to account for red noise by treating the residual as a mixture of Gaussians (68). Modification with other noise priors such as a student-t or Laplacian may also be beneficial.

**Parameter & hyper-parameter priors - $p(\theta)$:** The ability to place distributions over all parameters (and indeed all hyper-parameters) in Bayesian models is one of the great advantages of Bayesian inference. It requires selection of appropriate priors, which are updated to posterior distributions based on the data. We here detail the choice, and functional form, of these priors. Note that the prior over the parameters (weights) $w$ depends only upon a set of hyper-parameters which govern the scale (or precision) of the (multi-variate) Gaussian over $w$. We denote this set of hyper-parameters $\alpha = \{\alpha_k\}$. Further, we note that the prior over $\theta$ factorizes as:

$$p(\theta) = p(w|\alpha)p(\alpha)p(\beta). \tag{3.8}$$

The first two terms factorize as $p(w|\alpha) = \prod_k p(w_k|\alpha_k)$ and $p(\alpha) = \prod_k p(\alpha_k)$. Each of the terms in Equation 3.8 is described below.

$p(w|\alpha)$ - The model parameters (weights) $w$ are drawn from a zero-mean Gaussian prior. For element $w_k$ of $w$ this is represented as:

$$p(w_k|\alpha_k) = \left(\frac{\alpha_k}{2\pi}\right)^{1/2} \exp\left\{-\frac{\alpha_k}{2}w_k^2\right\}. \tag{3.9}$$

In this way, each quadrature amplitude (and additional fit coefficients: sea-level trend, MSS, etc.) is modeled by its own Gaussian distribution with inferred precision $\alpha_k$. This allows for the estimation of uncertainty over each weight $w_k$. This process is often referred to in the statistics literature as an *Automatic Relevance Determination* (ARD) prior (69), as it induces a shrinkage of $w_k$ values that do not contribute to model-data fit. Shrinkage refers to constraining the size and number of coefficients, thereby favoring smaller models unless the data require additional parameters.

$p(\alpha)$ - The weight precisions, $\alpha_k$, are each drawn from a Gamma (hyper)prior with initial mixing hyper-hyper-parameters $a_0 = 10^{-2}$ and $b_0 = 10^{-4}$ such that

$$p(\alpha_k) = \Gamma(\alpha_k; a_0, b_0). \tag{3.10}$$

The values $a_0$ and $b_0$ are selected to yield a vague prior. This simply means that the prior distribution is broad. Often Bayesian statistics will refer to this as an *uninformative prior.* This provides natural parameter shrinkage without imposing a strong pre-experimental assumption about parameter values. (70). Furthermore, the selection of a Gaussian prior for $w$ results in $\alpha$ being strictly positive, and thus a Gamma distribution is appropriate (71).

$p(\beta)$ - Finally, we model the precision (inverse variance) of the residual $\epsilon$ as a hyperparameter within the model. A Gamma prior is once again employed with mixing parameters $c_0 = 10^{-2}$ and $d_0 = 10^{-4}$ such that $p(\beta)$ is given by:

$$p(\beta) = \Gamma(\beta; c_0, d_0). \tag{3.11}$$

We note that $a_0, b_0, c_0, d_0$ are unitless as analysis can be run on normalized data, with final parameter estimates and predictions obtained by simply unnormalizing.

In the following section we briefly describe the principles behind our method of choice for performing inference: variational inference. This method allows us to update prior distributions over parameters of interest into the associated posterior distributions, conditioned on the observed data $Y$ given the assumed model $X$.

#### 3.2.2.1 Variational Inference

Our goal is to evaluate the posterior distribution $p(\theta|Y)$ for the harmonic model previously detailed. This evaluation requires us to infer distributions over the elements of $\theta$, namely the (linear) weights $w$, the associated weight precisions $\alpha$ and the precision, $\beta$, of the residual between model and target. Even for a linear model, this set of inferences is not analytically tractable. We could use Markov-Chain Monte Carlo (MCMC) sampling, which is an effective approach to approximating the required distribution. However, MCMC approaches can be computationally intensive and scale poorly as the number of parameters increases (68). This limitation becomes more salient when employing the spatially coherent harmonic analysis procedure we propose in Section 3.2.4. This is a consequence of the fact that the model simultaneously models the quadrature amplitudes for multiple sites which significantly increases the number of parameters. In our testing, the application of the Bayesian HA with MCMC sampling became computationally intractable beyond confined regions.

Instead, we advocate an approximate Bayes approach, referred to as *variational Bayes* (VBayes). Unlike, MCMC sampling, the VBayes framework offers a computationally tractable method of approximating the fully Bayesian inference of both parameters and hyper-parameters in the model. It achieves this by minimizing the Kullback-Leibler (KL) divergence between an approximate posterior $q(\theta|Y)$ and the true posterior distribution $p(\theta|Y)$. By imposing a functional form for $q(\theta|Y)$ which allows for computational tractability (72), variational inference reduces the optimization of probability distributions to a problem of sequentially estimating the hyper-parameters of the distributions. This retains the benefits of fully Bayesian modeling without the need for sample-based approaches. This section is meant to serve as an introduction to the key concepts (and limitations) of variational inference - a detailed derivation of variational inference is provided in (72) and a more detailed exposition of the Bayesian linear model, used here, can be found in the appendix of (71).

As variational inference is an iterative optimization procedure, models a are initialized using the maximum-likelihood (ML) solution given by

$$w_{ML} = (X^{\mathsf{T}}X)^{-1}X^{\mathsf{T}}Y, \tag{3.12}$$

where the $n^{th}$ row of $X = x_n$ and has the standard form given in Eq. 3.1 and $Y = [y_1, y_2, \dots, y_N]^T$. The ML solution is then used to initialize the residual precision hyper-parameter, $\beta$, such that:

$$\beta^{-1} = \frac{1}{N}\sum_{i=1}^{N}(y_i - w_{ML}^{\mathsf{T}}x_i)^2 \tag{3.13}$$

Variational inference operates by defining a strict lower-bound to the posterior (but intractable) likelihood, $p(\theta|Y)$. VBayes introduces an *approximate posterior* $q(\theta|Y)$. The functional form of this approximation is not arbitrary, but is in accordance with the prior over $\theta$, such that $q(\theta|Y)$ factorizes as:

$$q(\theta|Y) = q(w|Y)q(\alpha|Y)q(\beta|Y). \tag{3.14}$$

We then write the log of the data evidence, namely $\log p(Y)$, as the sum of two terms:

$$\log p(Y) = F\left(p(\theta|Y), q(\theta|Y)\right) + \mathrm{KL}\left(p(\theta|Y), q(\theta|Y)\right). \tag{3.15}$$

Equation 3.15 is the fundamental equation of variational inference. Understanding the two terms involved provides helpful intuition as to why this method is effective. The term $F(p, q)$ is known as the (negative) variational free-energy, or evidence lower bound (ELBO), and provides a *strict* lower-bound on the model evidence. Maximizing $F(p, q)$ therefore results in the approximate posterior being as close as possible to the true posterior. The second term denotes the Kullback-Leibler (KL) divergence between the approximate and true posteriors over $\theta$:

$$\mathrm{KL}(p(\theta|Y), q(\theta|Y)) = \int q(\theta|Y) \log\left(\frac{q(\theta|Y)}{p(\theta|Y)}\right) d\theta. \tag{3.16}$$

It can be shown that the KL term increases monotonically with the number of free parameters in $\theta$ and thus provides natural model shrinkage. Furthermore, the non-negativity of the KL divergence ensures the strictness of the variational bound.

Noting that the approximate posterior is factored (as per Equation 3.14), so $F(p,q)$ can be maximized by iteratively optimizing each of $q(\theta|Y), q(\alpha|Y), q(\beta|Y)$ separately. The update equations employed for this procedure create a series of sequential re-estimation equations with guaranteed convergence properties and are given in (73). Indeed, rather neatly, the ELBO of such a VBayes model is guaranteed to improve at each iteration. The more well-known expectation-maximization (EM) algorithm is a special case of the more general form (74), in which the distributions over the variables of interest are delta functions (and so yield only point-value estimates, rather than full distributions).

#### 3.2.2.2 Limitations and considerations when using VBayes

The VBayes approach is not without its limitations. First, we assume our model has conjugate priors. That is, the posterior $q(\theta, Y)$ will have the same form as the prior $p(\theta)$. In principle, this constitutes a strong pre-experimental assumption regarding the relationship between prior and posterior. We believe that our assumption of Gaussian distributed weights with Gamma distributed precisions is justified in the context of harmonic analysis after extensive analysis of constituent posteriors on synthetic and real data using a full Bayesian model in pymc-3 (a probabilistic programming language in Python) with MCMC sampling (75). Further the Gaussian and Gamma priors are also maximum entropy priors, and thus introduce minimal bias (76).

Second, VBayes (in the manner we use it) approximates joint distributions by products of marginals (72). This is a consequence of the fact that we assume the factorized marginal densities of Equation 3.14 are separable under the mean-field approximation (77). The resultant parameter estimation is generally accurate, but can miss out on local maxima of $p(\theta, Y)$. This is illustrated in Figure 3.1 Panel A. By capturing the maxima of the posterior distribution our model provides more robust parameter estimates, at the expense of the full shape of the marginals, and thus potential under-estimation of uncertainty. As will be described in Section 3.2.6.2, we integrate the modeled noise variance $\beta^{-1}$ into the final parameter level uncertainty estimates to reduce this under-estimation. In altimetric tidal analysis, we are primarily interested in understanding how tidal constituents co-vary. VBayes explicitly considers the covariance between parameters rather than assuming independence as in least-squares estimators. Given the interaction and correlation (multicollinearity) between these constituents, the joint distribution offers insight into the global behavior of all parameters, and thus more accurate parameter estimation. This is the opposite of least-squares, which effectively models marginal distributions by assuming parameter independence. In doing so, least-squares approaches may overlook the complex relationships that exist between tidal constituents, potentially compromising the robustness of parameter estimates and the associated standard error estimates.

Comparisons of the marginals and joint distributions using MCMC sampling could provide interesting insights into "cross-talk" and aliasing between constituents. This represents an interesting area of future work, as these interactions encompass both physical nonlinear hydrodynamic interactions and spurious signal processing errors.

### 3.2.3 Comparisons with other estimators

The rationale for using VBayes becomes clearer when constrasted with other estimators commonly applied in tidal analysis. Here we compare our formulation with the Bayesian formulations of several common estimators employed for tidal analysis: OLS, iteratively-reweighted least-squares (IRLS) (33), Lasso Regression (L1 regularization), and Ridge Regression (L2 regularization).

VBayes can be understood as a regularized weighted least-squares regression, albeit with more flexibility and sophistication due to the way these features are incorporated via the prior and the likelihood.

*Regularization:* Within a Bayesian context, regularization is achieved through the selection of priors over the inferred weights $w$. The Bayesian equivalent OLS and IRLS estimators do not impose a prior over these parameters, and they are thus unregularized. For Ridge regression, the prior is assumed to have the form of a zero-mean Gaussian such that $p(w) =$

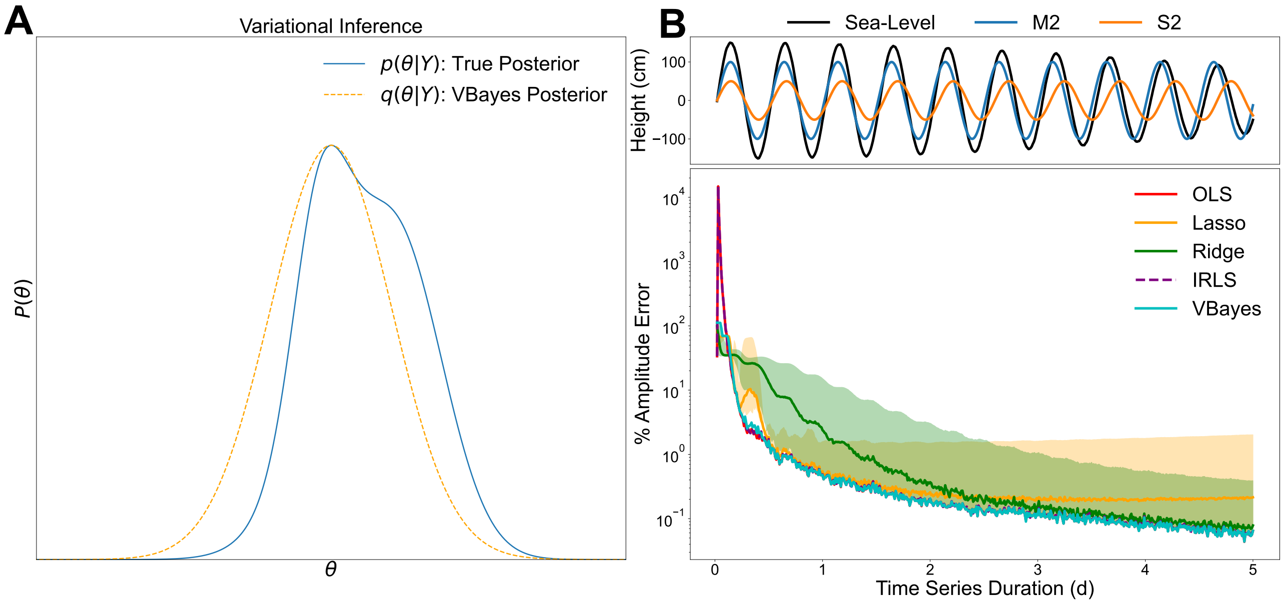

Figure 3.1: (**A**) Comparison between approximate VBayes posterior $q(\theta|Y)$ and true posterior $p(\theta|Y)$. (**B**) Comparisons of frequency resolution characteristics of different estimators. The synthetic signal is composed of two constituents M2 and S2 with amplitudes 100 cm and 50 cm respectively. The signal is contaminated with uniformly distributed noise in the frequency domain, as in (1), with approximate signal-to-noise ratio (SNR) 1000. Amplitude errors for each estimator are shown as the duration of the time-series used for estimation is increased. Intervals for Ridge and Lasso regression reflect the range of errors produced by varying $\lambda$ between 0.01, and 1. Note errors are shown on log-axis.

$\mathcal{N}(w \mid 0, \lambda^{-1} I)$. This prior encourages smaller parameters, but does not force them to be zero. In contrast, Lasso regression assumes the pdf to be a laplacian (double-exponential) where $p(w) = \prod_j \frac{\lambda}{2} \exp(-\lambda |w_j|)$. This prior has the effect of forcing less-relevant parameters to be zero.

The "strength" of these priors is controlled by the regularization coefficient $\lambda$, which must be selected beforehand. This can significantly impact the efficacy of the estimator as shown in Figure 3.1 Panel B. Similar to Ridge regression, our VBayes approach also imposes a Gaussian prior over the weights $w$. However, rather than have a predefined $\lambda$, we instead model the weight precisions $\alpha_k$ as an internal parameter within the model. This removes the need to select a regularization parameter beforehand, and instead lets the data inform the degree to which regularization is necessary. This characteristic is valuable for application to satellite altimetry, where the diverse set of conditions and massive datasets prevent manual or iterative cross-validation based selection of $\lambda$s.

As noted in (33) and can be seen in Figure 3.1 Panel B, IRLS and OLS estimators both exhibit a failure mode as the length-of-record is reduced. This failure mode leads to solutions containing significantly more energy than the original signal, and erroneous parameter estimation. Indeed, so too does the selection of an appropriate $\lambda$ for Lasso and Ridge. These deviations are notable as the theory of tidal resolution, and indeed super-resolution (1), has focused on unregularized least-squares. The impact of the VBayes priors is readily seen, as the VBayes does not have the characteristic failure mode of OLS and IRLS, while also avoiding bias introduced by the fixed regularization when sufficient data exist.

*Weighting:* While regularization encourages less complex models by promoting sparsity or smaller parameter values, problems arise in the presence of outliers, which violate the assumption of Gaussian residuals. Even with regularization, these outliers can significantly bias parameter estimates. The impact of outliers can be mitigated by *down-weighting* their influence. In a Bayesian context, this is achieved through the likelihood function, which can adjust for outlier effects in different ways.

As described above, the VBayes likelihood is given by

$$p(Y|X,\theta) = p(Y|X,w,\beta) = \left(\frac{\beta}{2\pi}\right)^{N/2} \exp\left\{-\frac{\beta}{2}\sum_{i=0}^{N}(y_i - w^\intercal x_i)^2\right\},$$

where $\beta$ is inferred alongside $w$. The inverse precision, $\beta$, weights the influence of the individual data points, modeling the uncertainty based on the observed noise level. Hence, by fitting $\beta$ within the model, the second term in the likelihood (Equation 3.6) weights the observations based on the weighted residuals ($\beta^{-1}$) such that observations with smaller residuals contribute more to the model fit and *vice versa*. We note that OLS, Ridge, and Lasso regressions all assume a fixed $\beta$ estimated directly from the residual, and thus do not down-weight outliers in this way.

For the case of IRLS, three primary differences exist. First, rather than model $\beta$ within the model, a tuning parameter $\tau$ is used, which scales the influence of outliers on the model (e.g. small $\tau$ down-weights outliers more and *vice versa*). Just as with Lasso and Ridge regression, a weakness of this approach is the need to specify this parameter beforehand which can lead to over or under-fitting. Conversely, by modeling $\beta$ explicitly, the equivalent parameter tuning in VBayes is accomplished implicitly through posterior inference. The second difference in IRLS is that the scaled residual is fed into a 'weighting function', such as Huber's function or Tukey's bi-square. This is equivalent to modifying the assumed functional form of the residual in the likelihood to a more heavy tailed distribution. For these weighting functions, the most similar associated pdfs would be either a Laplace or Student's-t distribution. Third, IRLS introduces an additional weighting matrix, $W$, which explicitly weights each data-point. In a Bayesian context, this would be equivalent to introducing an additional set of latent weights into the model which weights each measurement within the likelihood function (78). In this way, while our VBayes approach down-weights the impact of outliers without requiring a tuning parameter, the more explicit down-weighting within the IRLS approach of leffler2009enhancing can be seen as more robust to outliers, if an appropriate tuning parameter can be selected. However, while this can increase model robustness, our testing found IRLS to struggle on the SWOT Cal/Val data. We attribute this to the extreme sparsity of the data. To increase the robustness of our approach to outliers, we introduce an automated outlier detection procedure in Section 3.2.5 using the intrinsic uncertainty estimation with our method.

### 3.2.4 Spatially Coherent Harmonic Analysis

Inspection of early SWOT data confirms our physical intuition that the oceanic response is spatially coherent. This fact, while interesting, cannot be fully exploited using single point measurements such as tide gauges. Empirical tidal models from satellite altimetry already capitalize on this, even if indirectly, through binning of measurements into fixed grids (79; 19). Similarly, several works have attempted to use nearby gauges to further constrain solutions through interpolation (80). As will be discussed, SWOT data has revealed conventional binning, averaging, and interpolation procedures can mask short-scale features in the derived tides. Swath based altimetry provides high-resolution spatial measurements whose spatial information can be exploited directly if an appropriate method can be developed. Here we develop a spatially coherent inductive bias which can be imposed on the model. We note that an inductive bias differs from a prior, in that an inductive bias imposes an assumption about the relationship between data points (here we assume a spatial relationship), whereas a prior refers to our prior belief about the distribution our parameters come from.

Spatial coherence simply implies that there is information relevant to our point of interest in the surrounding points. Mathematically we can define this as a single convolution between the central point $p_{0,0}$ and the surrounding points $p_{j,k}$. Recall that our general harmonic analysis model can be defined as the linear model $Y = w^{\intercal}X + \epsilon$. Once again, $Y$ is the set of all measurements, $X$ is an $M \times N$ matrix of basis functions where $M$ is the number of measurements and $N = 2n+2$ with $n$ equal to the number of constituents, and $w$ is the set of derived weights. To allow information to propagate to the central point we modify this procedure to simultaneously regress the neighboring points by stacking the measurements, and associated weight and design matrices. This imposes the assumption that the sea-surface at neighboring points $p_{j,k}$ is given by the sum of the sea-level at the central point $p_{0,0}$ and an additional offset.

Measurements and design matrices are padded with zeros for neighboring points such that after dropping zero coefficients the sea-level at $p_{j,k}$ is predicted by

$$Y_{j,k} = X_{0,0}w_{0,0} + \rho \cdot d \cdot X_{j,k}w_{j,k} \tag{3.17}$$

where $d$ is the normalized distance between $p_{0,0}$ and $p_{j,k}$ taken to be $(0,1)$ for uniformly spaced data, $\rho$ is the probability that the observations $Y_{j,k}$ are correlated with $Y_{0,0}$. This is obtained by taking the Fischer transform of the Pearson's correlation coefficient $r$ between $Y_{0,0}$ and $Y_{j,k}$, such that $z' = 0.5[\ln(1+r) - \ln(1-r)]$, and $\rho = 1 - \Phi(z')$ with $\Phi = \mathcal{N}(0,1)$. This weighting imposes a simple physical assumption; the more similar the time-series is between two points, the closer their harmonic constituents should be. Hence, in the limit where both points are identical, $Y_{j,k} = X_{0,0}w_{0,0} + 0 \cdot 1 \cdot X_{j,k}w_{j,k}$. This formulation is equivalent to a first-order spatial autoregressive process, or SAR (81). By simultaneously inferring weightings across points based on the central point, the central point can be better resolved. While this procedure down-weights the entries of the design matrix for the neighboring points, the ratio of parameters to data-points is actually higher than in the single-point case. One can do better if one imposes more restrictive assumptions regarding the relationship between neighboring points.

For closely spaced points (the present approach has been tested up to 4 km spacing), we modify the above procedure based on the fact that the relationship between the quadrature amplitudes of neighboring points is approximately linear. This assumption allows us to describe the relationship between the center point and its neighbors using only two lines. The following procedure is designed for the structured grid-data produced for the SWOT L2 data products but is easily adapted to unstructured data. The lines are defined by

$$Y_{j,k} = \rho \cdot (d_x \cdot X_{j,0}w_{j,0} + d_y \cdot X_{0,k}w_{0,k}) + X_{0,0}w_{0,0} \tag{3.18}$$

where $d_x$ and $d_y$ are the normalized distances between points along the horizontal and vertical direction respectively, taken to be $(-1, 0, 1)$ for convenience. For SWOT these can be thought of as representing the distance along the along-track and cross-track directions respectively. A 3-d visualization of what this looks like for a set of points is provided in Appendix A. Using four neighboring points yields an increase in the effective number of measurements to parameters by nearly a factor of two (1.6 to 2.9). It is possible to incorporate more than four neighboring points on a non-uniform grid by weighting the distances appropriately. The results presented in this paper only utilize points separated by up to 4 km (e.g. only consider neighboring SWOT pixels). As will be shown in Section 3.3.2.1, this spatial coherence can significantly improve the accuracy of constituent estimation when using both sparse data and data with noise artifacts. The decision to use only nearest neighbors in this work reflects our objective of incorporating spatial coherence without imposing unrealistic assumptions about the spatial distribution of amplitudes. While amplitudes may smoothly vary (in a linear sense) for much longer distances, the quadrature amplitudes are not necessarily linear and so care should be taken before applying this procedure to coarser data. Further validation is needed to determine where the breakdown point lies for the linearity assumption.

This procedure can be leveraged using any conventional estimator, including OLS and its variants. While we believe the VBayes procedure offers several intrinsic advantages, conventional OLS augmented with spatial coherence worked well in low-noise regimes. These results are given in the supplementary materials (Figure S3).

### 3.2.5 Bayesian Pruning

Due to the presence of erroneous or contaminated measurements in coastal regions, statistical procedures for outlier removal typically fail to remove values contaminated by other non-stationary artifacts (e.g. storm surge, roll errors, satellite maneuvers, etc.). While the VBayes approach is generally more robust to such artifacts than OLS, their presence will still cause biasing of the derived constituents. Implicit within the VBayes framework is uncertainty estimation over both the parameters and the data. This feature can be leveraged to flag and even remove points which exhibit behavior inconsistent with our model's expectation (e.g. stationary tidal signal with a linear trend and mean sea-level). The procedure is remarkably simple. For each measurement, the model computes a predictive distribution with mean $\mu_i$ and standard deviation $\sigma_i$. The computed standard deviation is composed of the contribution $\sigma_i = \sigma_p + \sigma_{\text{noise}}(i)$ where $\sigma_p$ is the sum of the standard deviations from the derived constituents and $\sigma_{\text{noise}}(i)$ is the contribution to the standard deviation from the $i^{th}$ measurement. Measurements which sit greater than $3\sigma_i$ away from the mean are removed. The selection of this threshold was empirically determined using synthetic data. Testing of this procedure is provided in Section 3.3.2.2. We note this algorithmic pruning procedure echos the arguments provided in (33) while removing the need for parameter tuning which creates complications for unsupervised application to satellite measurements.

### 3.2.6 Constituent Selection

Arguably the most important consideration in any harmonic analysis is the choice of which constituents to include. These decisions have major implications for the accuracy and bias of the derived harmonics (10). Indeed, these decisions are made more salient when attempting to derive tidal constituents from sparsely sampled data, as is often the case with early stage SWOT data. Several methods have been proposed to determine which constituents can be obtained from a given reference series. The most well-known, and most often used, is the so-called Rayleigh criterion. Rayleigh defines the minimum uniformly sampled time-series duration ($T$) in order to separate two constituents $\omega_1$ and $\omega_2$ as

$$\left(\frac{T}{\omega_1 - \omega_2}\right) / R_{min} \geq 1, \tag{3.19}$$

where $R_{min}$ is typically taken to be 1. Numerous works have shown that this criterion is often overly restrictive and that in practice one can resolve additional constituents depending on the signal-to-noise ratio (SNR) (1). As shown in Figure 3.1 Panel B, regularized estimators exhibit different frequency resolution characteristics from their unregularized counterparts, and can often perform better particularly in the presence of noise. As such, in practice we find even modified constituent selection approaches which consider SNR to be insufficient. Further discussion on this is provided in Section 6.5. Rayleigh is often considered in conjunction with the Nyquist sampling theorem which states that for a given frequency $\omega$ to be resolvable, the associated signal must be sampled with sampling frequency $f \geq 2\omega$ (37). Due to the fact that the tidal signal has a sparse Fourier basis, accurate constituent estimations can be obtained using sampling frequencies far below the Nyquist frequency (38) as is confirmed by the success of empirical tidal models from historical altimetry (19).

We would like to emphasize that the selection of constituents goes beyond avoiding aliasing. Even for constituents which are resolvable according to Rayleigh, simple testing using OLS HA on synthetic data demonstrates that inclusion or exclusion of a given constituent can degrade/improve the accuracy of other constituents significantly. This is a consequence of how co-linearity, or cross-talk (10), interacts with different estimators. Simple comparison of different estimators yields different results. Hence, we argue that to maximize the effectiveness of a given analysis, the choice of constituents must be done in conjunction with the estimator employed for analysis.

To address this we propose a data-driven approach which can efficiently explore and produce optimal combinations of constituents for any uniform or non-uniform sampling period. This approach is used to obtain a robust global optimal constituent solution for the SWOT Cal/Val orbit.

#### 3.2.6.1 Evolutionary Optimization for Constituent Selection

Due to the massive data volume from the KaRIn, an iterative approach to constituent selection is computationally infeasible. However, based on the known sampling period ($\approx 0.99349$ days for Cal/Val), we can determine an optimal solution for our variational Bayesian harmonic analysis using synthetic data generated from the TICON-3 dataset (52). The dataset provides a diverse set of tidal regimes containing various levels of non-linearity and is thus suitable for our purposes when searching for robust solutions. In order to efficiently explore potential solutions we utilize a genetic algorithm (GA) to perform evolutionary optimization. We define a candidate solution as a binary string with each digit corresponding to a constituent. A value of 1 means the corresponding constituent is used in the analysis, and a value of 0 means it is excluded. We restrict the set of candidate constituents to the 40 contained within the TICON-3 dataset (52). While this set of constituents accounts for a majority of global tidal variance, the lack of consideration of additional constituents and potential influences of their co-linearity is a notable limitation.

Evolutionary optimization is performed using a multi-objective loss function and is tasked with simultaneously minimizing the RMS error of the M2 constituent, while maximizing the number of constituents with median percentage errors below 10%. The choice of median percentage RMS error is commensurate with our objective to produce a robust model, though we note that the optimal solution for median percentage RMS error is also the optimal solution for the mean RMS error. Interestingly, we also found that the optimal solution for M2 often coincided with the optimal solutions of several other constituents. By definition, the best parameter estimates will occur when the estimator is the least biased. This occurs when the residual is independently distributed with equal variance. Hence, the optimal set of constituents will be the set which most closely satisfies this assumption. The TICON-3 dataset is used to generate 2569 pure tidal signals which are contaminated with moderate white Gaussian noise (100 SNR) and nonstationary noise (one event with signal-to-surge ratio of 0.1). Details for how signals are contaminated with white and nonstationary noise is provided in Sections 3.3.1.1 and 3.3.1.2 respectively. Additionally, the GA implementation can be found in the VTide GitHub repository[1]. For each candidate solution, 300 stations are selected at random from the TICON-3

---

[1] `https://github.com/thomasmonahan/VTide`

dataset, time-series are then generated based on the SWOT Cal/Val sampling period and noise is applied. The evolutionary optimization is run for 100 generations (30,000 models) with the objective of obtaining a representative set of candidate solutions. This dataset provides interesting insights into the relationship between individual and sets of constituents and the Cal/Val sampling period.

A visualization of the top solutions (Figure A.1) indicates strong positive and negative "lines" for different constituents. However, noise exists depending on which constituent we choose to optimize for. To account for this, we select the top 20 combination of constituents for the primary constituents M2, N2, S2 and sum the total number of occurrences for each constituent using both median percentage RMS error and average RMS error. The resulting sums are divided by the total number of possible occurrences to yield a probability of inclusion. We select an empirically determined threshold for inclusion of 80% based on experimentation using both synthetic and real SWOT data. For the real data, inappropriate selection of constituents leads to visually "fragmented" phase estimates where interacting constituents will bias others. This interrupts the spatial coherence which is expected and observed when appropriate constituents are selected. It was found that the smoothest solutions were produced by the 80% threshold.

The global optimal solution found is comprised of M2, N2, S2, O1, M4, and MN4. We find that this deviates slightly from the optimal solution for semi-diurnal regions which also includes M6 and M8. In some instances, such as the Bristol channel, we've found that many additional minor constituents such as M10, MU2, 2N2, NU2, and MU2 can reliably be obtained, however, they are not necessarily reliable outside of macro-tidal regimes. The dominant diurnal constituent K1 was found to significantly degrade the accuracy of the analysis likely due to close proximity between the K1 period and the SWOT Cal/Val sampling period (1.0027 vs. 0.9934 days). While the evolutionary procedure is imperfect, it provides an alternative to practitioner trial and error, the imperfections of Rayleigh, and provides us with a reliable solution for unsupervised applications. We note, however, that while this makes sense for SWOT with the fixed sampling period, it is computationally impractical for standard harmonic analysis. Under these conditions, the ARD priors, in conjunction with the superior uncertainty estimation, could aid in constituent selection but is not considered further due to the sparsity of the SWOT data.

#### 3.2.6.2 Bayesian Constituent Uncertainty Estimation

It remains to be said how the quadrature amplitudes $(\mu_A, \mu_B)$ and associated amplitude variances $(\sigma_A^2, \sigma_B^2)$ can be related to the final uncertainty estimate. The following procedure closely relates to the linearized analysis in (32).

First, we compute the equivalent amplitude and phase $C = \sqrt{\mu_A^2 + \mu_B^2}$ and $\phi = \arctan(\mu_B, \mu_A)$. Next, we assume the contribution of the noise variance $\sigma_{\text{noise}}^2 = \beta^{-1}$ (aleatoric uncertainty) to the given constituent's uncertainty is proportional to the relative magnitude of the constituent $\upsilon = C / \sum_i^{2N} w_i$. It follows that the total variance associated with each component is given by $\sigma_{A_{tot}}^2 = \sigma_A^2 + \upsilon \cdot \sigma_{\text{noise}}^2$, and $\sigma_{B_{tot}}^2 = \sigma_B^2 + \upsilon \cdot \sigma_{\text{noise}}^2$. Through a laborious derivation, the amplitude variance can be expressed as

$$\sigma_C^2 = \left(\frac{A^2}{A^2 + B^2}\right)\sigma_A^2 + \left(\frac{B^2}{A^2 + B^2}\right)\sigma_B^2, \tag{3.20}$$

with the phase variance given by

$$\sigma_\theta^2 = \left(\frac{B^2}{(A^2 + B^2)^2}\right)\sigma_A^2 + \left(\frac{A^2}{(A^2 + B^2)^2}\right)\sigma_B^2. \tag{3.21}$$

These values can then be used to generate RMS uncertainty estimates using the RMS formula given in Section 3.6). Robust uncertainty estimates in amplitude and phase are achieved by incorporating both epistemic and aleatoric uncertainties. This yields dynamic and locally accurate uncertainty estimates as shown in Section 3.4.1.

## 3.3 Simulated Results

### 3.3.1 Comparisons Between OLS and VBayes for SWOT Cal/Val Orbit

Prior to assessing the impact of the additional modifications outlined in Sections 3.2.4 and 3.2.5, it is worth first comparing the behavior of the OLS and VBayes methods under different conditions. To do this, we once again utilize the TICON-3 dataset (52). While this cannot represent all global tides, the data provides a comprehensive distribution of tidal regimes (Semi-Diurnal, Mixed, Diurnal), tidal ranges (micro, meso, macro), and varying degrees of shallow-water effects. Synthetic reference series are generated for each tide gauge at the SWOT Cal/Val sampling period (0.99349 days) unless explicitly stated otherwise. This sampling does not represent crossover-points, which will have twice as many measurements. The two methods are then compared with variable types of noise contamination and sampling periods in the following sections. Models are evaluated on the Root Mean Squared error (RMS) (defined in Section 3.6) of the primary constituents in the "optimal" genetic algorithm (GA) solution ($M2$, $N2$, $S2$, $O1, MN4, M4$). The root-sum-of-squares (RSS) of all included constituents is also presented. The definition for the RSS is given in 3.6. The present analysis only compares OLS and VBayes as OLS is the relied upon method for present empirical tidal analysis of altimetry data (19). While this work focuses on conventional harmonic analysis, the VBayes estimator is similarly applicable to the response method (5; 40). As described in Section 3.2.3, OLS is not the only estimator employed for tidal analysis. In our testing, we found IRLS to perform significantly worse than both the VBayes and standard OLS on data with SWOT sampling periods. Due to the near 50% increase in RSS error we have not included it in the subsequent analysis. Additionally, given the need to specify a $\lambda$ parameter beforehand, Ridge and Lasso regressions are also not considered further.

#### 3.3.1.1 White Noise

To test the two method's resilience to uncorrelated and normally distributed noise, the simulated reference series are contaminated with white Gaussian noise (WGN) at signal-to-noise ratios (SNR) ranging from 0-100. For each SNR, trials are repeated 100 times for all simulated series with the mean and median values for the 100 trials presented in Figure 3.2 Panel(s) B. It can clearly be seen that for high SNR ratios, the OLS and VBayes are approximately equal. However, as the noise floor is increased, there is a pronounced divergence at approximately 40 SNR in which the OLS starts to perform significantly worse while the VBayes remains relatively bounded. This behavior is consistent across the individual constituents and the RSS error. Interestingly, while the WGN satisfies the normality assumption of the error term by the OLS, the OLS still breaks down as the WGN is increased. This is a consequence of the fact that the OLS estimator's "breakdown point", defined as the minimum amount of contamination needed to cause an estimator to take on erroneous values (82), is proportional to the inverse of the sample size (83). Conversely, the compact support supplied by the priors in the Bayesian framework yield an estimator which in principle never breaks down (82). The degree of this support, however, will of-course depend on the strength of the priors.

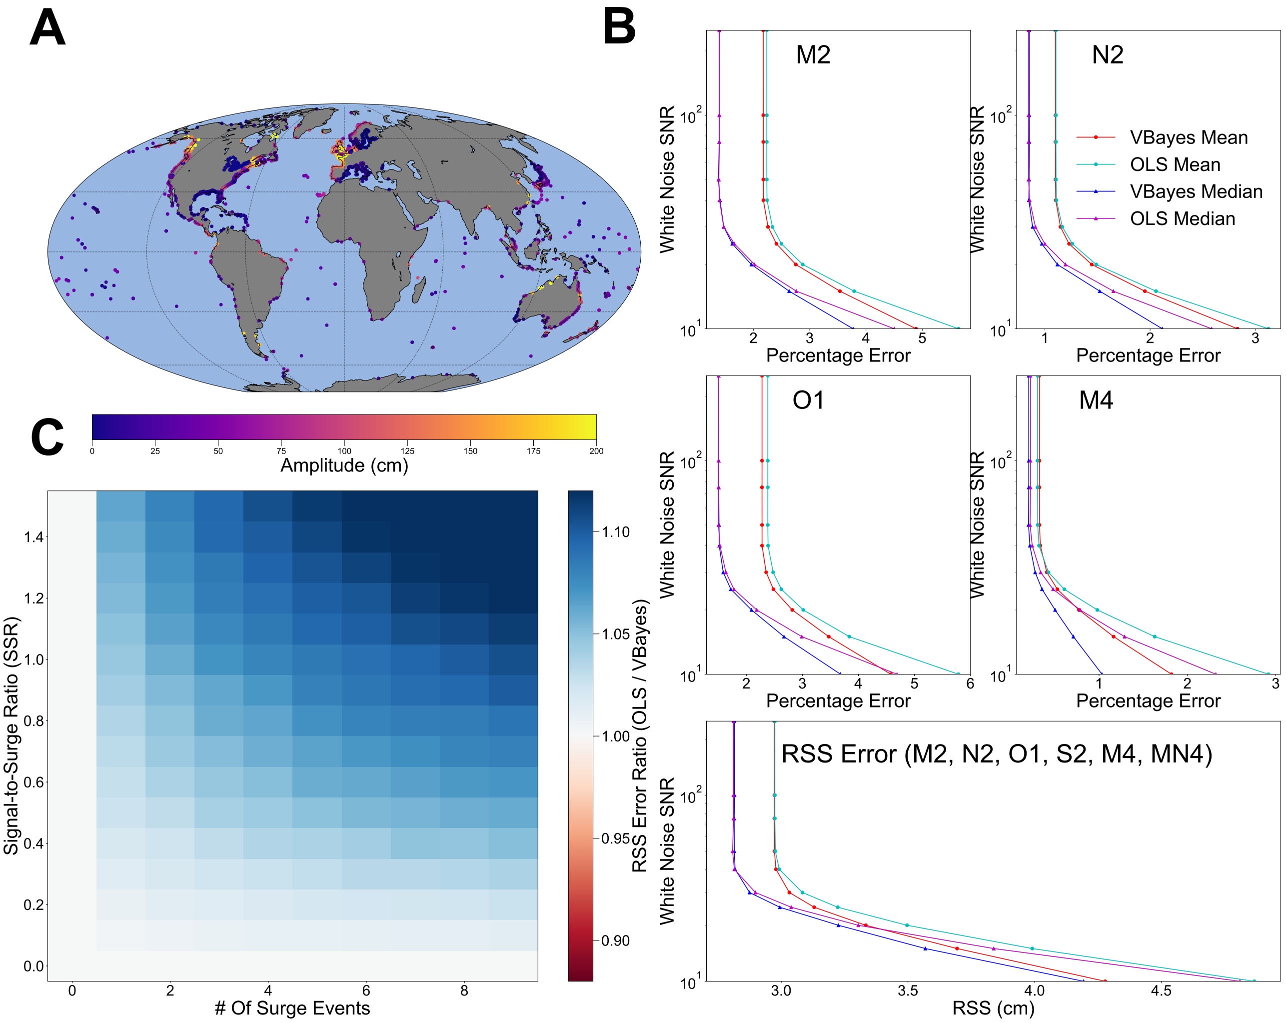

Figure 3.2: **Comparisons of OLS and VBayes on synthetic TICON-3 dataset when contaminated with different types of noise.** Panel A shows the distribution of TICON-3 tide gauges used to generate the synthetic reference series employed for testing. Panel B compares OLS and VBayes when contaminated with increasing amounts of white gaussian noise. The mean and median values for the VBayes are given in red and blue. The mean and median values for the OLS are given in cyan and magenta. Panel C compares the ratio of OLS to VBayes RSS values at the associated signal-to-surge ratio and number of events (Blue = VBayes better, Red = OLS Better, White = Equal). For both Panels B and C the RSS is computed for the constituents (M2, N2, S2, O1, M4, MN4).

#### 3.3.1.2 Non-Stationary Noise

To simulate non-stationary artifacts we contaminate the simulated tidal signals with several synthetic surge events. We utilize the formula outlined in Cartwright's work on surges (16) such that a synthetic surge is given by $\zeta_S = A \cdot (1 + \cos(2\pi t/l))$ where $6 \leq l \leq 96$ (hours) is the duration of the surge, defined from $0 \leq t \leq l$ and $0.1 \leq A \leq 1.5$ is the "surge-to-signal ratio" (SSR) of the surge defined by the ratio of the surge magnitude $\zeta_S$ to one half of the tidal range. Tests are repeated by varying both the number of surges $N$ occurring within the 90-day Cal/Val period and the associated SSR. The range of SSR values and the frequency of events is based on empirical analysis of the residuals from the UK operational surge model (84). In reality, both extremes (negligible noise, or extreme noise) are unrealistic. Hence, the objective of this case study is to show the range of possible scenarios to allow researchers to make informed decisions about the performance of each estimator. Surges are randomly added to the signal without overlap. For each SSR and number of occurrences, tests are averaged over the full TICON-3 dataset (2569 signals) and the mean error ratios defined as $RMS_{OLS}/RMS_{VBayes}$ are plotted in heatmaps in Figure 3.2 Panel C. Signals are contaminated with a baseline 100 dB of $WGN$ to prevent exact inversion by OLS. Similar to the behavior shown in Section 3.3.1.1, Figure 3.2 Panel C shows that for minimal non-stationary contamination the two methods are statistically equivalent. However, as the number of surge events or surge intensity is even slightly increased, the two methods diverge with the VBayes yielding superior performance. This behavior conforms exactly to what is theoretically predicted by the two estimators when applied to data contaminated by non-normally distributed noise. While the present analysis has used synthetic surges to simulate non-stationary contamination, similar results would be expected in the presence of other non-stationary artifacts (e.g. roll errors, signal dropout, etc.).

#### 3.3.1.3 Sampling Period

Due to the global disparities in revisit time for locations sampled in the Science phase of the SWOT mission, and to demonstrate the salience of our method for application on historical altimetry data, we compare the performance of the two methods as the sampling period is changed. Three-year synthetic signals are generated with increasing sampling periods from 1 hour (Standard Tide Gauge Sampling Period) to 21 days (SWOT Science orbit repeat period).

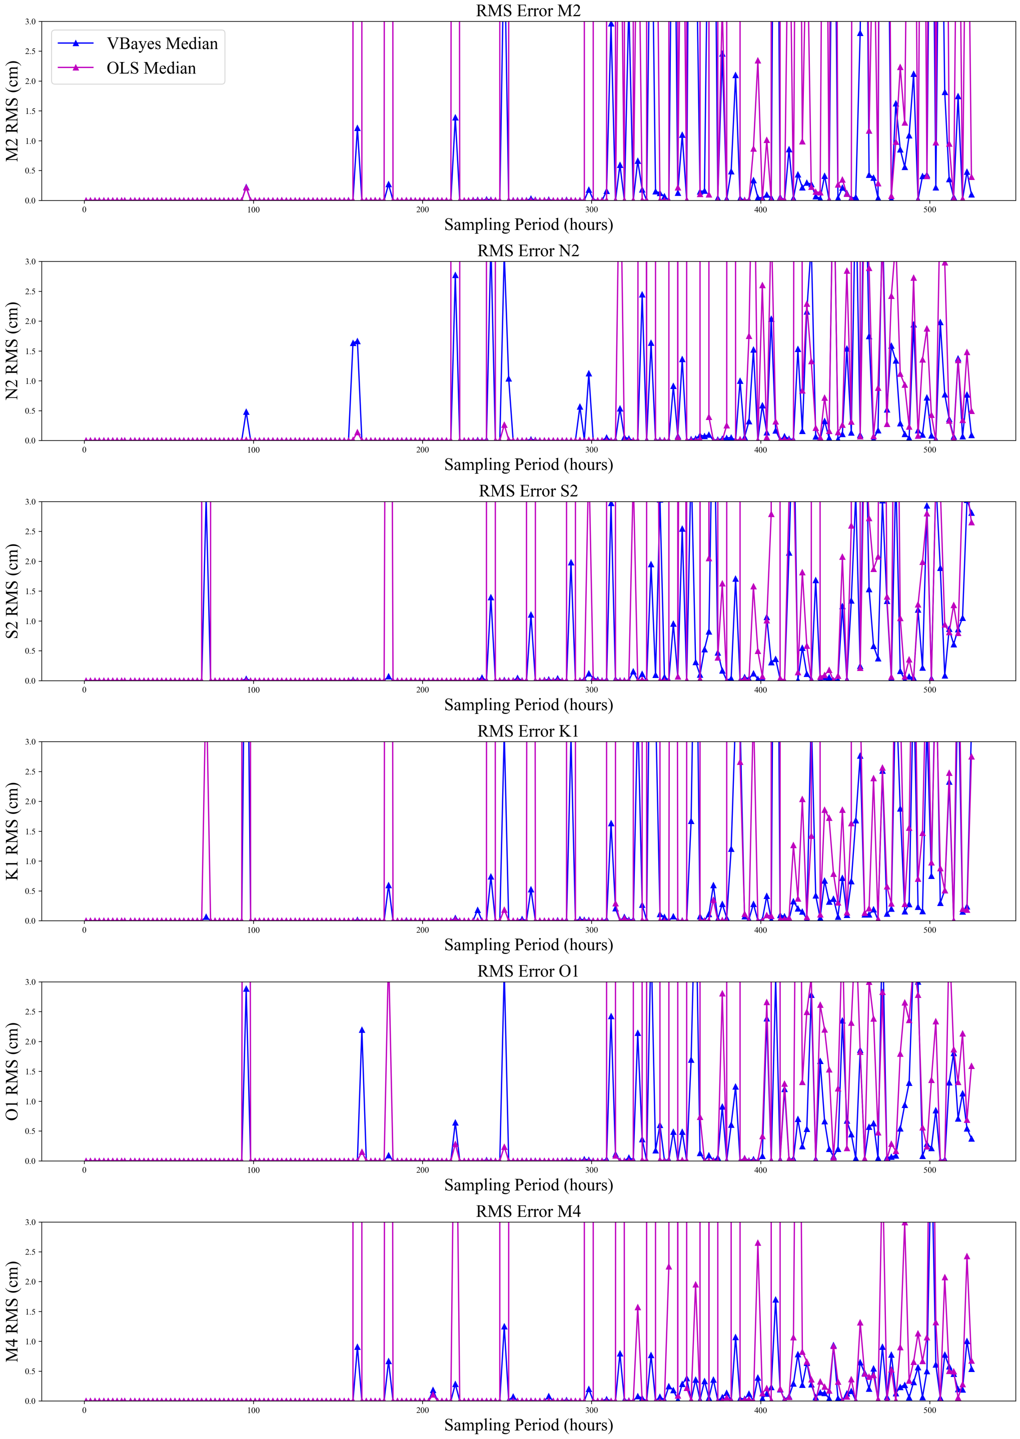

Figure 3.3: **Comparisons of RMS Error for OLS and VBayes at different sampling rates.** Median RMS error is computed for the entire TICON-3 dataset based on a 3-year reference series spanning the SWOT Science Orbit sampled at the given sampling rate.

During the SWOT Science orbit, locations will be sampled multiple times within a given cycle, thus the sampling period will vary geographically. Once again a baseline 100 dB of $WGN$ is added to prevent exact inversion. VBayes and OLS are run on the signals using all 40 harmonics in the TICON-3 dataset in the analysis. Figure 3.3 shows the median RMS error of several constituents and the associated RSS of each estimator and the associated constituent. Compared to the clear improvements in robustness shown in Sections 3.3.1.1 and 3.3.1.2, the performance characteristics of each estimator are far more complex when varying the sampling period. The duality of the time-frequency domain is, of course, a contributing factor – interactions between the sampling rate and constituent frequencies can induce aliasing and distortion effects. It can be seen that the two estimators behave differently for different constituents (e.g. sometimes OLS exhibits “aliasing” behavior while the VBayes does not, and *vice versa*). In general, we observe that the VBayes is more robust to true aliasing. This “robustness” comprises two aspects: constraining the derived values of aliased constituents, and ensuring the errors these present do not propagate to other non-aliased constituents. Both of these aspects are supported by the ARD priors over the weights. This result is particularly promising for sun synchronous missions such as Envisat in which problems arrise if solar constituents are included in an analysis. Indeed, even when combining these orbits with non-sun-synchronous missions errors can be introduced using ordinary least-squares. Simulated results for several constituents (including S2) when sampled at the Envisat sampling period (35 days) are provided in 3.3. The reduced biasing of the other constituents is a consequence of the regularization from the Bayesian ARD priors (69) which reduce the influence of parameters that do not contribute to model-data fit. However, further study is needed to understand this behavior. We anticipate these characteristics could be further exploited to utilize both estimators in tandem to offset periods when one outperforms the other.

### 3.3.2 Simulated Testing of Spatial Coherence and Bayesian Pruning

Now that the disparities between conventional OLS and VBayes have been established, we now look to validate the spatially coherent harmonic analysis as well as the Bayesian pruning procedures proposed in Sections 3.2.4 and 3.2.5 respectively.

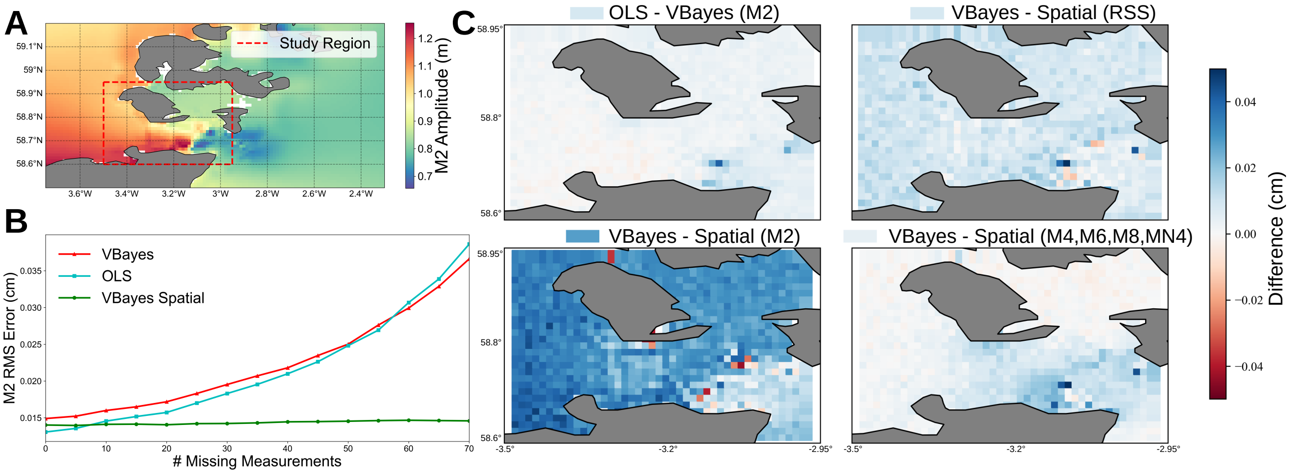

Figure 3.4: **Validation of the VBayes spatial coherence.** Panel A shows the simulated M2 amplitudes in the Pentland Firth and the study region used in Panels B and C. Panel B compares OLS, VBayes, and the spatially coherent VBayes as datapoints are randomly dropped (originally 102 points). Panel C compares the average difference in RMS error between the given models when contaminated with spatially incoherent noise (Red = first model better, White = same, Blue = second model better).

#### 3.3.2.1 Spatial Coherence

To validate the spatially coherent harmonic analysis procedure, we extract harmonics from a year long 2-D depth averaged simulation of the tides in the Pentland Firth, UK. The model is forced with constituents M2, K1, K2, MU2, N2, NU2, O1, and S2 at the ocean boundary and has been tested and validated against *in-situ* measurements in (85). The tidal heights in the Pentland Firth are largely semi-diurnal and are characterized by elevated shallow water effects induced by islands throughout the channel and some of the fastest tidal currents globally. The presence of simulated tidal stream turbines in the center of the channel amplifies these effects and produces relatively fast changes in tidal amplitudes due to the head drop as energy is extracted. The data offers a challenging stress test of the spatial coherence in complex coastal environments. The non-uniform grid-points are binned into a 2 km × 2 km grid to simulate the SWOT ocean data products and we restrict our study region to the center of the channel which is characterized by the greatest spatial variability. A visualization of the M2 amplitudes for the full simulation, and the 3041 node study region is shown in Figure 3.4 Panel A. We test the proposed method under several conditions; random data dropout, spatially incoherent noise, and spatially coherent noise.

Early SWOT data has revealed the importance of model robustness under data-dropout which can arise from several sources including tidal flats, satellite maneuvers, and other roll errors. While these data may be reprocessed in the future, presently in the Bristol Channel for example ≈ 10% of the Unsmoothed 250 m × 250 m measurement locations contain fewer than 70 measurements out of the 102 possible with the maximum number of measurements being 84 before outlier removal. Figure 3.4 Panel B compares the standard VBayes, OLS, and spatially coherent VBayes M2 RMS error as the number of randomly dropped measurements is increased. For this example, the spatial VBayes is only given access to a single neighboring point to capture "worst-case" performance. At least one neighbor is available in almost all cases for the SWOT data products. As before, 100 dB of WGN is added to prevent exact inversion by OLS and no non-stationary noise is added. As the number of missing measurements is increased the VBayes and OLS M2 RMS error appears to increase quadratically, with the OLS RMS error increasing by over 200% when the number of measurements available is reduced from 102 to 62 (0-40 missing data points in Figure 3.4). Conversely, the spatially coherent VBayes appears to increase linearly as measurements are dropped, with a relative increase of 18% using 32 data points (72 missing data points in Figure 3.4) from the central point. This is comparable to the OLS using more than double the number of measurements (72). In testing, the inclusion of additional neighbors further increased absolute model performance and robustness, though the improvements are less significant than those from the addition of the first point. As expected, both the spatially coherent and standard VBayes models perform marginally worse than OLS when no data dropout occurs. Under these conditions the assumptions of the OLS holds and it is thus equivalent to the maximum-likelihood estimator. It therefore outperforms any non-linear unbiased estimator by definition (62). The Spatial coherence can also be implemented for the OLS. An analogous plot to Figure 3.3.2.1 Panel B is given in Appendix A (Figure A.3) and further confirms the benefits garnered by the Spatial coherence. The OLS is more sensitive to deviations from the linearity assumption and contamination from non stationary noise and thus care should be taken if using the spatially coherent OLS.

Next, we evaluate the proposed method when applied to signals with both spatially incoherent noise and moderate dropout. Synthetic signals for the 3041 nodes in the study region are contaminated with 100 dB of WGN, and are subjected to five spatially incoherent non-stationary events ranging from 0 – 1.0 SSR. Signals also receive random dropout ranging from

0 – 50 points leaving 77 measurements on average. These contaminants attempt to emulate the conditions observed in early SWOT data in coastal regions which are characterized by much higher noise and erroneous measurements. The four plots in Figure 3.4 Panel C compare the average OLS, VBayes, and VBayes spatial differences by location and constituents over all trials (367,961 total simulations). As described in Section 3.2.6.1, it was found that in predominantly semi-diurnal regions it is also possible to obtain accurate estimates of the minor constituents M6 and M8. These were also included in the analysis to test the impact of spatial coherence on minor constituent resolution. As shown previously, the standard VBayes garners a slight improvement in RSS over the OLS under moderate non-stationary contamination. Comparisons between the standard VBayes and the Spatially coherent VBayes show a significant improvement in RSS under the same conditions. The source of these improvements can be seen in the bottom two panels. The largest improvement is in the dominant M2 constituent shown in the bottom middle panel. The Spatial VBayes on average demonstrates a relative improvement in M2 RMS of 44% and 40% over the OLS and standard VBayes respectively. Interestingly, estimation of other major constituents (N2, S2, and O1) remain relatively unchanged, though, critically they do not degrade. The other area of improvement produced by the spatial coherence is the resolution of the nonlinear and compound constituents (M4, M6, M8, and MN4). The difference in RSS for these constituents between the standard and spatially coherent VBayes is shown in the bottom right panel. Whereas the lowest RMS improvements for the M2 are found near the tidal turbine line in the center of the channel, the Spatial VBayes demonstrates the largest improvements in minor constituent RSS in this region. Under significant noise contamination, these minor constituents can be lost beneath the noise floor. However, the spatial coherence helps to preserve the coherent information beneath the spatially incoherent noise. These improvements are significant as minor and compound constituents are critical to tidal correction accuracy in coastal regions and also drive many important coastal processes such as tidal asymmetry.

Finally, we evaluate the proposed method under spatially coherent noise. The same noise contamination utilized for the tests in Figure 3.4 Panel C is applied to the synthetic signals. However, this time non-stationary / white noise contamination are applied to all signals simultaneously. That is, if non-stationary noise is added to a point, it is added to all neighboring

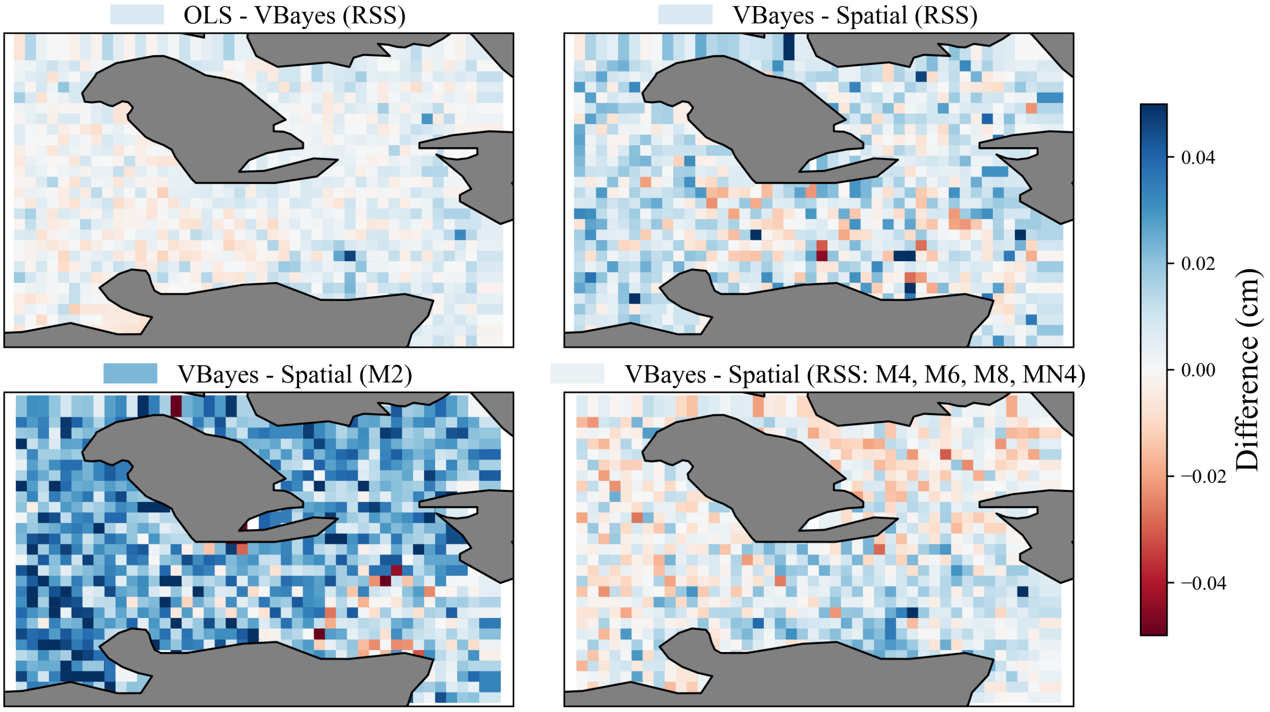

Figure 3.5: **Comparisons of model performance under spatially coherent noise.** Panels compare the average difference in RMS error between the given models when contaminated with spatially incoherent noise (Red = first model better, White = same, Blue = second model better). The average difference is given next to the title of each plot.

points at the corresponding time and with intensity drawn from a normal distribution given by $\mathcal{N}(SSR, SSR/10)$. Equivalent plots for the spatially coherent noise are provided in Figure 3.5. Unsurprisingly, the spatial coherence performs worse when faced with this type of contamination, though it still yields superior performance to both the standard VBayes and OLS. As will be shown in the following section, the Bayesian pruning procedure can be used to offset these deficiencies.

#### 3.3.2.2 Bayesian Pruning

In order to validate the Bayesian pruning procedure we must address two relevant questions. First, can our model correctly identify erroneous points? Second, due to the already limited amount of available Cal/Val data, does removing these datapoints improve performance? To assess these points we contaminate the data with non-stationary noise exactly as done in Section 3.3.1.2. Then, by applying the procedure outlined in Section 3.2.5, erroneous points are iteratively flagged by the model and removed. It can be seen that as outliers are pruned away, the VBayes prediction improves relative to the true signal and has greater confidence in the

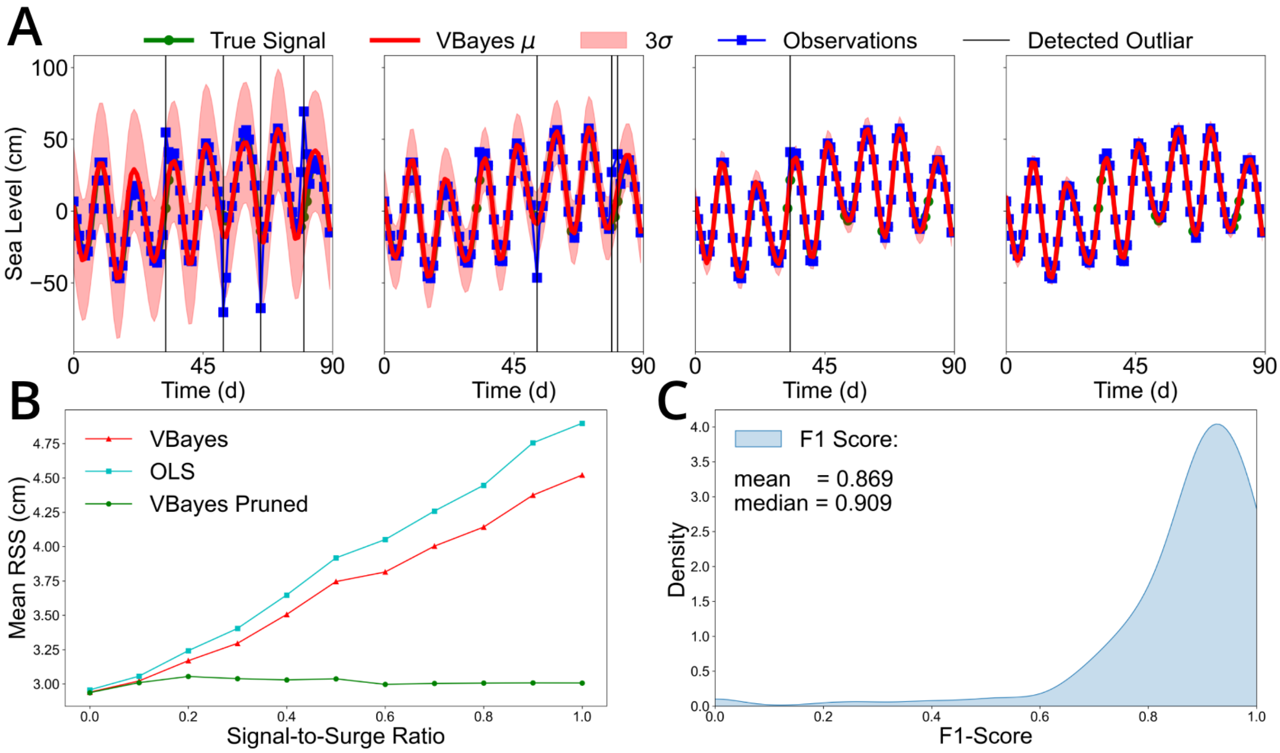

Figure 3.6: **Evaluation of the pruning procedure outlined in Section 3.2.5.** Panel A shows the iterative pruning procedure, with outliers being identified, removed, and the procedure repeated in each subsequent panel. Panel B compares the mean RSS error for the standard VBayes, OLS, and VBayes with pruning across M2, N2, S2, O1, M4, MN4, M6 and M8 when applied to TICON-3 synthetic signals contaminated with 100 dB WGN, SSR ratio ranging from 0-1.0, and number of events set to 5. A kernel density estimate of the distribution of F1-scores across all trials, defined by the binary classification of a point as being contaminated or not, is shown in Panel C.

derived constituents (Figure 3.6 Panel(s) A). Pruning can be viewed as a simple binary classification problem. Hence, we assess the performance of our method using the F1-score which provides a balanced measure of precision versus recall for positive values (outliers). The formula to compute the F1-score is given in Section 6.6. A kernel density estimate of the distribution of F1-scores for the tests is shown in Figure 3.6 Panel C. The average F1-score of 0.860 and median score of 0.909 confirm the model's efficacy when classifying erroneous points. By selecting our threshold to be $3\sigma$, we prioritize minimizing false positives. Due to the fact that we are only including a subset of possible constituents in our analysis, minimizing false positives is crucial to not throwing away other tidal variability. This is reflected by the false positive rate of 1.451% and the true positive rate of 56.55%. It should be noted that the true-positive rate is likely underestimated because the synthetic surges range over multiple days, and even minor surge contributions are considered positive.

Now that we have verified that our model can accurately detect outliers we must evaluate how the derived constituents are impacted by their removal. Figure 3.6 Panel B compares the mean RSS of the OLS, standard VBayes, and VBayes with pruning when contaminated by 5 non-stationary surge events ranging for $0.0 \leq SSR \leq 1.0$ and WGN at 100 dB. Whereas the standard VBayes and OLS errors appear to increase monotonically with the added non-stationary noise, the pruned VBayes RSS error increases only marginally from the noise-free case. It can therefore be concluded that the pruning procedure is both robust and effective at removing such contaminants (e.g. surges, land contamination, roll errors, etc.). Notably, and intentionally, this procedure will likely not remove systematic contaminants acting across longer time-scales, such as river outflow, and thus will not prevent biasing from these types of noise.

## 3.4 SWOT Data

Now that the VBayes methodology has been extensively validated using synthetic data, we apply our method to real SWOT data. As of writing this (May 2024), there are several versions of the SWOT Cal/Val dataset of varying resolutions and processing levels. Here, we focus on the L2 2 km gridded ocean product (Processing level B) as this is intended to be the primary vehicle for scientific research and contains full geophysical corrections. We focus our analysis on the European continental shelf due to the abundance of tide gauges present within the Cal/Val orbit and the fact that two separate passes cross-over this region (003 and 016). We extract all available data from passes 003 and 016. Here, all empirical analysis is conducted on the sea-level anomaly (SLA) unless explicitly stated otherwise. The SLA is derived by applying all default geophysical corrections including the CNES mean sea surface, the FES2014 tide model, and roll error correction. While tides have been partially removed using the FES2014 model, we here look to test our model's ability to identify and remove the residual variability. The validation of the methodology is carried out in two parts. First, we compare the VBayes derived harmonics directly with the 22 gauges in the TICON-3 dataset. Next, we measure the associated SLA variance reduction for all locations in passes 003 and 016. The variance reduction is defined as the difference between the variance of the SLA before and after applying the tidal correction. This metric is built on the idea that if the model is correctly removing systematic tidal signals, the residual variability should be lower. All results are presented with uncertainty estimates and with examples of how these can be used to determine the confidence in a given correction.

### 3.4.1 Gauge Comparisons

Tide gauge validation is carried out on the subset of the TICON-3 gauges contained within passes 003 and 016 of the SWOT Cal/Val orbit. Due to processing errors in the Bristol Channel, UK we do not make use of the tide gauges located in this region for validation. Gauges are selected based on several criteria. First, we locate all gauges which fall within 4 km of any SWOT KaRIn L2 2 km measurement location from passes 003 and 016. This yields an initial set of 45 gauges. For a location to be included we enforce that there must be at least 40 measurements present. We further restrict the set of measurements to those which have at least one neighbor in order to validate the spatial VBayes. Next, seven duplicate gauges were identified and removed to avoid over-representation from a given site. In total, 22 gauges were used for validation and are shown in Figure 3.7 Panel A. It is worth noting from the outset that due to the large tides in this region and the presence of several estuaries, the improvements garnered by empirical approaches will be larger than those globally.

During testing, it was found that certain constituents were more reliably estimated through residual analysis of the SLA data than direct analysis of the uncorrected data, and *vice versa.* We note the "residual" in residual analysis, refers to the fact that the analysis is carried out

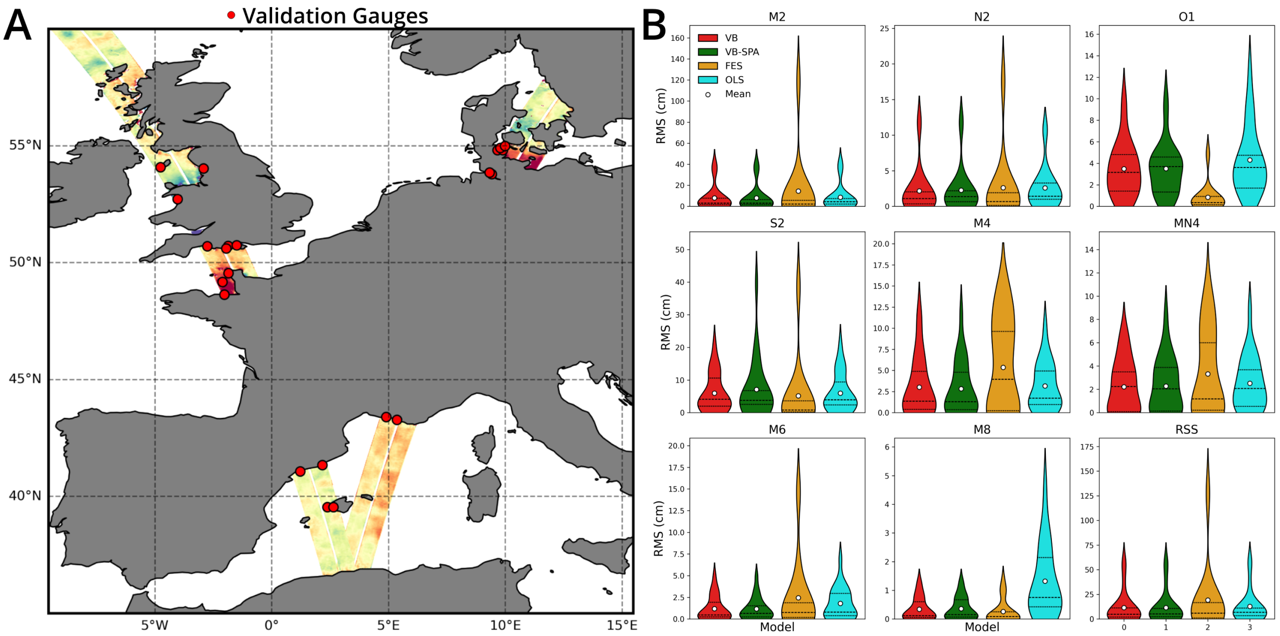

Figure 3.7: **Tide gauge validation.** Panel A shows example SSHA data for passes 003 and 016 over the European Continental shelf. The 22 TICON-3 tide gauges contained within these swaths are shown in red. Panel B compares the distribution of RMS across the 8 constituents identified in Section 3.2.6.1 and the total RSS between the VBayes (Red), VBayes Spatial (Green), and FES2014 (Orange) and OLS (Cyan) models. The mean values are denoted by the white circles. Solid horizontal lines represent the 25% and 75% quantiles. The dashed line represents the median.

on the SLA which is the residual signal after applying the included tidal corrections. These differences appear to be systematic and can thus be exploited. Similar optimization has been done within the context of the inference of minor tides (86). For the results shown in Figure 3.7 the M2 and N2 constituents are obtained via residual analysis with the remainder being obtained through direct analysis of the sea-surface height (SLA - FES Corrections). For the selected gauges it can be seen in Figure 3.7 Panel B that the Spatial VBayes yields a 46% reduction in average M2 RMS relative to FES. This large improvement is explained by the fact that the distribution of FES RMS errors exhibits heavy tails due to poor performance in the Elbe estuary, GER. Estuaries remain an open challenge for global tidal modeling. Large and statistically significant mean improvements produced by the Spatial VBayes over FES can also be found for the N2, M4, MN4, and M6 constituents with relative improvements in RMS error of $46\%, 14\%, 48\%$, and $32\%$. As discovered in the synthetic testing, the empirical methods struggle to resolve the S2 and O1 constituents. However, inclusion of S2 and O1 significantly improves the quality of the other derived constituents, particularly M2 and M4. These errors

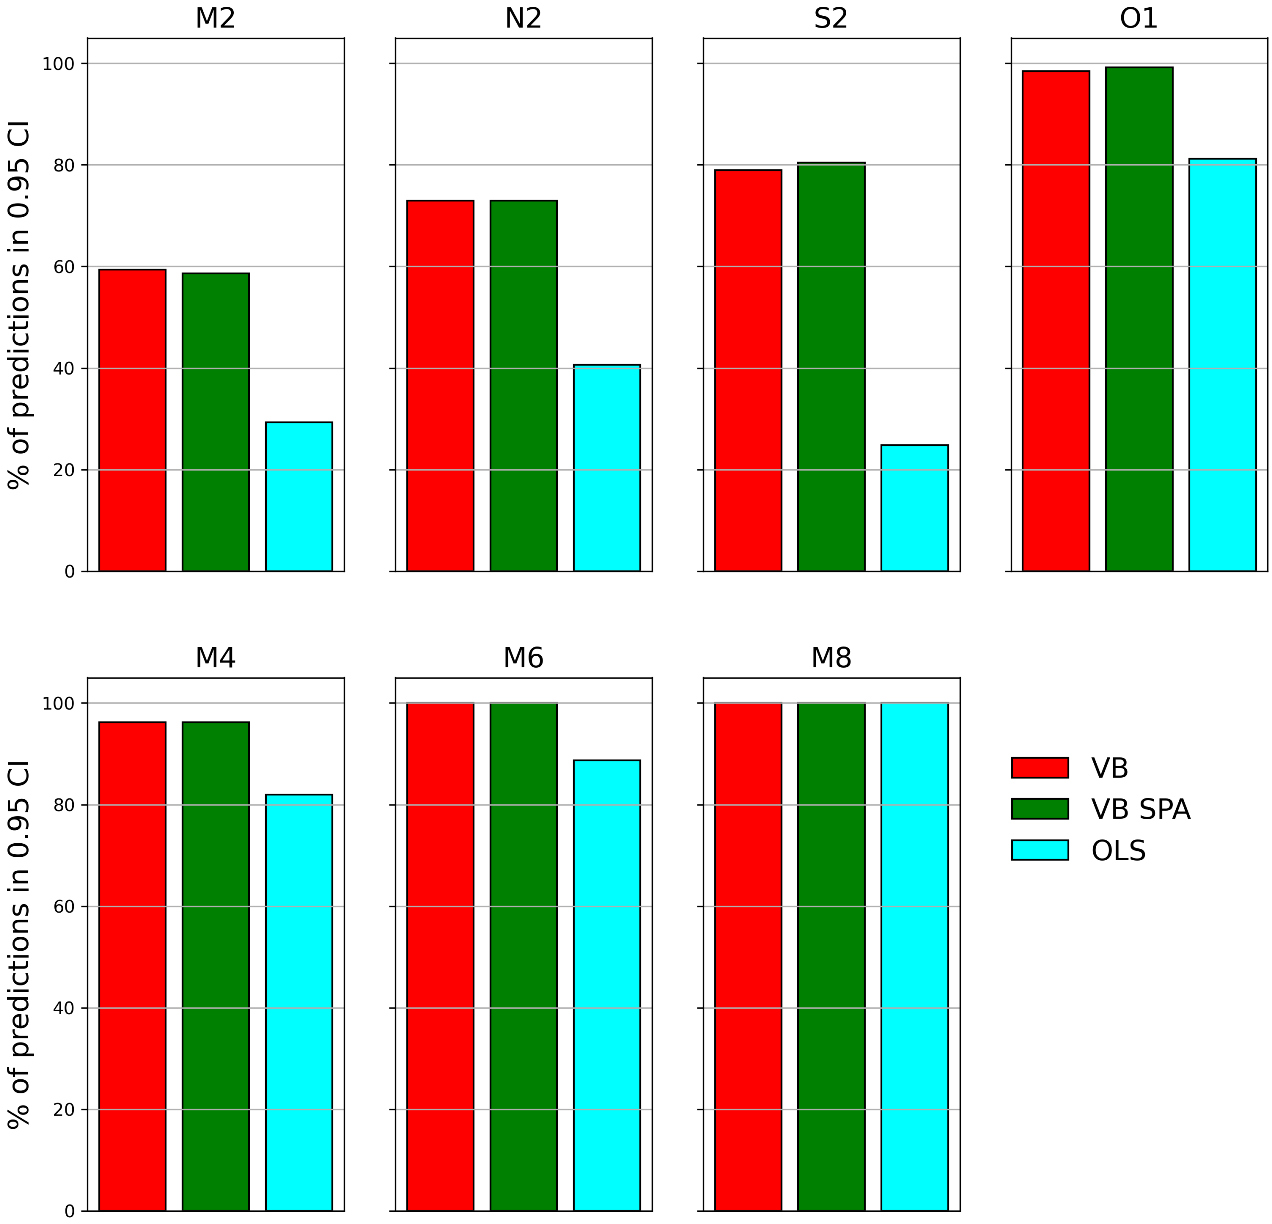

Figure 3.8: **Comparisons of VBayes and OLS uncertainty estimates for Gauge Validation.** The percentage of predictions which fall within the computed 0.95 Confidence Intervals (CI) are shown. A total of 133 SWOT measurements were used (all available measurements within 4 km of the gauges used).

are nicely reflected by the VBayes uncertainty estimates, with 81% and 99% of the spatial VBayes estimates falling within the 0.95 CIs for S2 and O1 respectively. Further comparisons of the coverage of the equivalent VBayes and OLS CIs for the gauge validation are shown in Figure 3.8. It can be seen that OLS significantly underestimates the uncertainty of most constituents, particularly S2. The spatial VBayes marginally improves CI coverage relative to the standard VBayes. While all methods can be seen to underestimate the uncertainty for M2 and N2, both the VBayes CIs provide a more than 30% increase in coverage relative to OLS for these constituents.

Finally, FES yields a marginal improvement in M8 resolution. However, this change is not statistically significant. With these results in mind, and in keeping with our empirical correction philosophy, our final model utilizes the Spatial VBayes estimated M2, N2, M4, MN4, and M6, and combines them with the remaining FES constituents. The resulting RSS estimates are shown in the bottom right panel of Figure 3.7. Through supplementing the constituents derived by the Spatial VBayes approach, the average FES RSS error in this region can be reduced on average by 40% with a median reduction of 12%. This improvement in median RSS indicates the Spatial VBayes using Cal/Val SWOT data demonstrates systematic improvements in constituent resolution over FES in addition to just improving erroneous regions. We anticipate these results, and the number of additional constituents that can be recovered by Spatial VBayes will increase as more SWOT data becomes available. It should also be noted that many locations contained only one neighboring point for the Spatial VBayes and thus, improved performance is expected where additional measurement locations are available. Finally, we believe the statistical approach to estimate residual tidal variability presented in the next section can provide an unsupervised framework to only apply the empirical corrections where we are confident the present corrections are insufficient.

Inspection of the coastal tides derived from SWOT data, particularly using the L2 Unsmoothed 250 m and Pixel-cloud products, has revealed that tides exhibit rapid spatial variability in and around coastlines (57). This variability constitutes a non-negligible source of error when validating against tide-gauges. Classically, models are validated by either selecting the closest point to a gauge or performing linear interpolation. These approaches have been necessary but can constitute a large source of uncertainty in complex coastal regions (57). Our

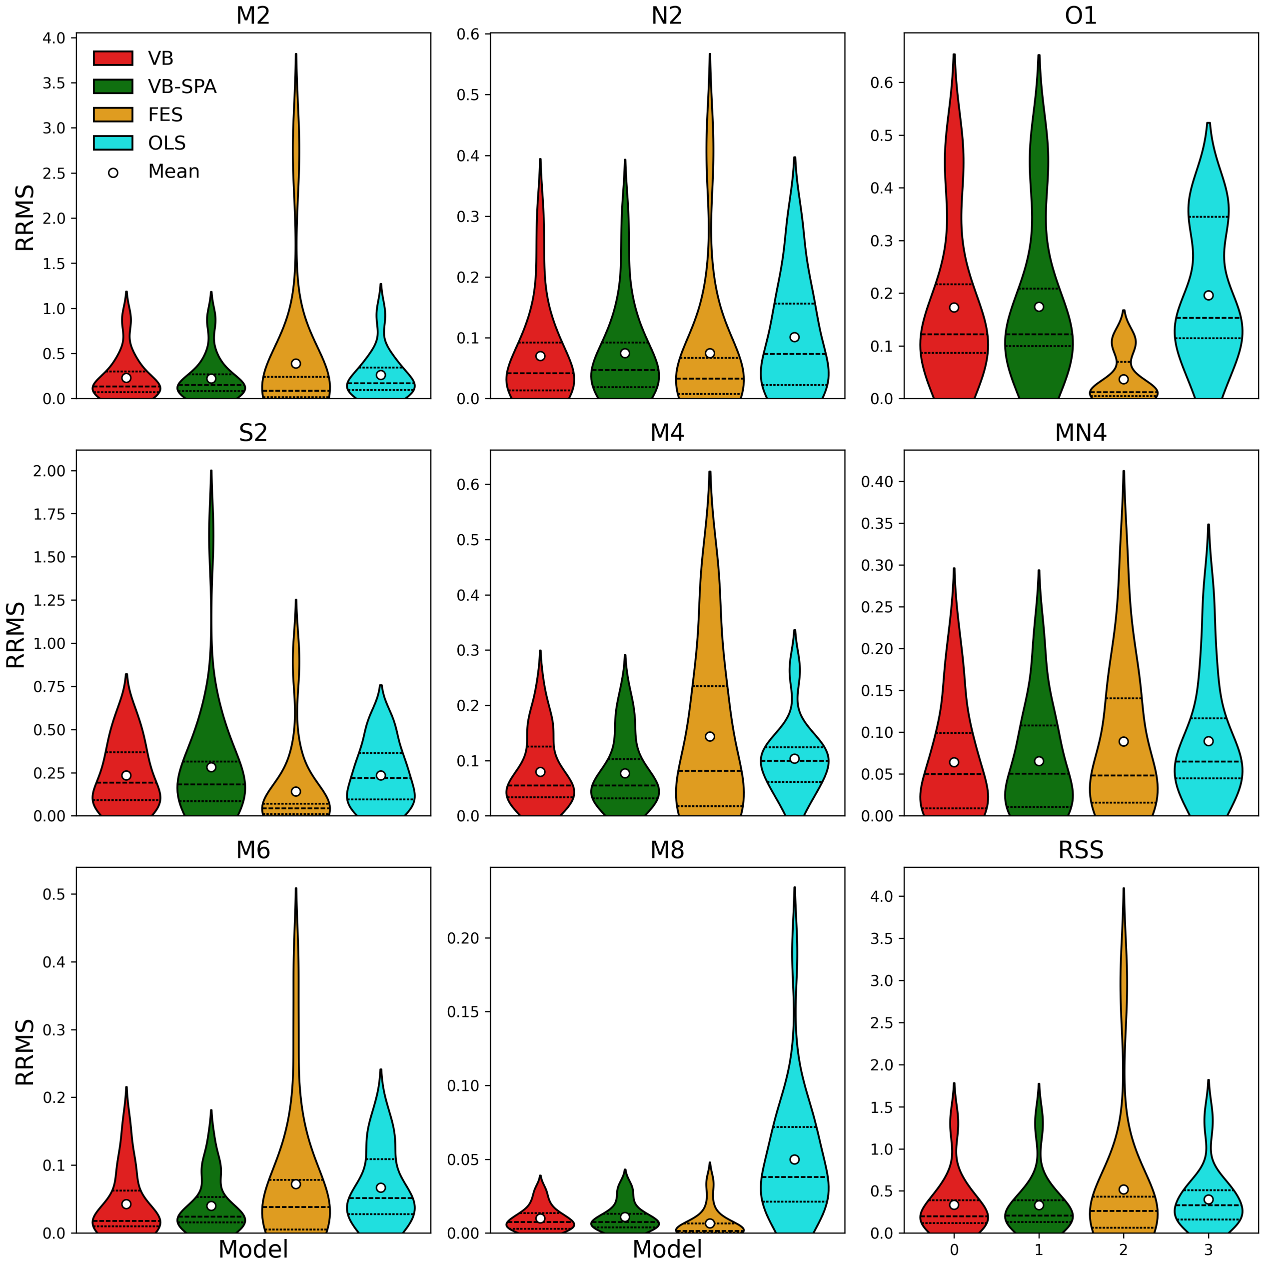

Figure 3.9: **Comparisons of model relative RMS (RRMS) error for tide gauge validation.**

own analysis identified several gauges exhibiting significantly different behavior from the observed SWOT data. The identification and documentation of such gauges will be important for future tidal model validation. Quantifying the magnitude of these errors is important as they can yield biased estimates of model performance. This is particularly salient for models which assimilate gauges directly. To control for the relative error this introduces we propose a relative RMS metric ($RRMS_k$) which weights the RMS error based on the mean absolute difference between the tide gauge harmonic signal and the observed SWOT data. This takes the form

$$\text{RRMS}_k = \frac{\sqrt{[(A_{\text{model}} \sin(\omega_k t) + B_{\text{model}} \cos(\omega_k t)) - (A_{\text{true}} \sin(\omega_k t) + B_{\text{true}} \cos(\omega_k t))]^2}}{\sum_{t=0}^{N} |y_{obs}(t) - y_{gauge}(t)|} \tag{3.22}$$

where the numerator is the conventional constituent RMS given in (23), and the denominator is the mean absolute error (MAE) between the SWOT measurements and the tide gauge harmonic prediction. We favor MAE over mean squared error (MSE), as MSE places greater weight on outliers. This metric quantifies the error of the tide model relative to the "uncertainty" introduced between gauges and adjacent measurements. Figure 3.9 presents the same analysis as Figure 3.7, but uses the RRMS metric. While the Spatial VBayes continues to yield superior mean RRMS for M2, N2, M4, and M6, the shape of the RRMS distributions changes significantly. The RRMS distributions for the empirical models exhibit significantly less kurtosis relative to the RMS. Due to the fact that the RRMS down-weights errors associated with larger differences between the SWOT and gauge measurements, this reduction in kurtosis for the empirical models indicates that many of the erroneous RMS values can actually be attributed to discrepancies between the gauge location and the SWOT measurement location. By accounting for these discrepancies, we can conclude that any remaining large errors in the RRMS values are due to systematic modeling errors in whichever model is being assessed. This is certainly the case for the dominant FES outliers in the Elbe estuary which remain when using RRMS. While we do not believe this metric should replace the standard RMS, it can add an additional dimension with which to understand the systematic errors intrinsic to the tide gauge validation of global tide models.

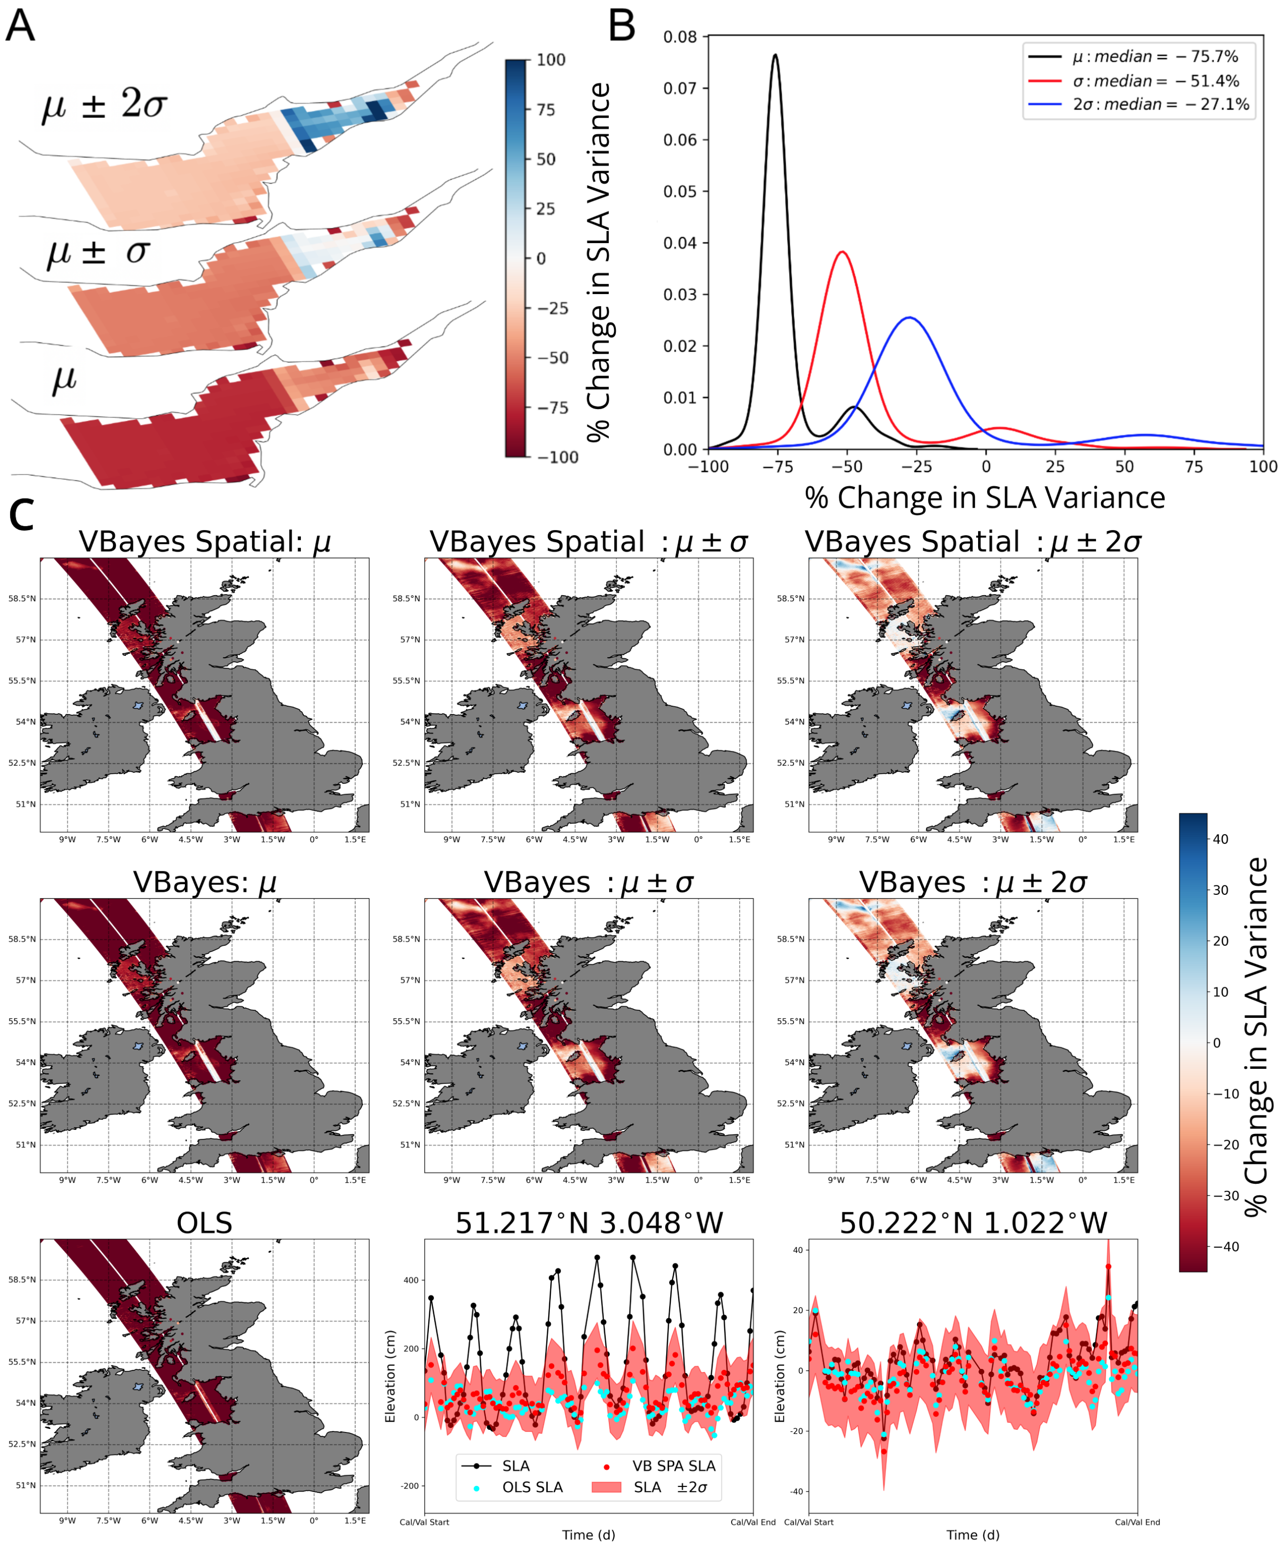

Figure 3.10: **Comparisons of SLA variance change after applying empirical corrections under different confidence levels.** Panel A shows the Spatial VBayes percent variance change in the Bristol Channel, UK when considered with uncertainty. Panel B shows the associated kernel density estimates of the distributions of these changes. Panel C compares the Spatial VBayes, standard VBayes, and OLS when used to perform residual analysis on data from pass 003 around the UK at different confidence levels. The bottom two panels show representative examples of the tides from the Bristol Channel, UK, and the English Channel respectively. Blue = increase in variance, Red = reduction in variance. Variance difference is given by $Var_{diff} = 100 * (\frac{Var_{FES} - Var_{model} - N\sigma_{model}}{Var_{FES}})$ where $N$ is the number of standard deviations of uncertainty to include.

### 3.4.2 SLA Variance Reduction

While *in situ* tide gauge validation allows for direct accuracy assessment of the derived constituents, point estimates cannot effectively quantify the spatial characteristics of model performance. Estimation of such errors is particularly salient in complex coastal regions where SWOT data has revealed standard linear interpolation to be deficient. Furthermore, several of the gauges employed for the previous analysis were assimilated into FES2014b and can thus yield biased estimates of global accuracy (23). We look to approximate the spatial error characteristics of FES by applying the proposed methodology to correct the residual SLA. We estimate the SLA variance difference according to $Var_{diff} = 100 * (\frac{Var_{FES} - Var_{model} - N\sigma_{model}}{Var_{FES}})$ where $N$ is the number of standard deviations of uncertainty to include. Note, OLS estimates do not include uncertainty. Inspection of the SLA data reveals several large errors likely due to satellite manoeuvres and roll errors which lead to high SLA variance estimates. These errors are removed by filtering the initial SWOT data for values exceeding $3\sigma$ beyond the sample mean. The three methods are then applied to the filtered SLA with the results shown in Figure 3.10.

Panels A and B visualize the percentage SLA variance reduction in the Bristol Channel, UK, using the spatial VBayes under different uncertainty thresholds. Due to the massive tidal range and severe nonlinearity in this region global tide models typically perform poorly (57). A visualization of an SLA signal in this region can be seen in the second column in the bottom row of Panel C. Clearly, considerable tidal variability remains, as is evidenced by the reduced SLA variance after applying the OLS and the standard and Spatial VBayes. It can also be seen that a majority of this variability exists beyond the $\pm 2\sigma$ confidence intervals of the Spatial VBayes. Thus, we can say with 95% confidence that the original SLA which sits beyond this region can be attributed to residual tidal variability. Higher up in the Severn estuary it can be seen that the model is less confident. This is a combination of increased nonlinearity, river outflow, interactions with large sand deposits, and sensor dropouts which lead to large increases in aleatoric uncertainty (e.g. uncertainty from the noise intrinsic to the data itself). Even for confidence levels of 68% and 95% respectively, the Spatial VBayes still yields median SLA reductions of 51.4% and 27.1% respectively (Figure 3.10 Panel B). Hence, the inclusion of the Spatial VBayes corrections to these data can significantly improve the quality of the derived

SLA for other studies. At the time of writing this, issues with the processing of the new SWOT low-rate, version B, data products have resulted in large gaps in coverage in the Bristol Channel. As such, the results shown in the Bristol Channel study makes use of the version A data-products and associated processing. While our testing showed little to no change in SLA variance reduction between the A and B data products for the portion of the channel which was present in both data products, we stress that caution should be taken when interpreting these results.

| Method | $\mu$ | $\mu \pm \sigma$ | $\mu \pm 2\sigma$ |
|---|---|---|---|
| VBayes Spatial | 43.3% | 25.2% | 7.2% |
| VBayes | 42.7% | 24.5% | 6.4% |
| OLS | 53.4% | – | – |

Table 3.1: **Average estimate of residual tidal variance for SWOT Cal/Val orbit passes 003 and 016 under different confidence levels.** $\mu \pm \sigma = 68\%$, $\mu \pm 2\sigma = 95\%$.

Figure 3.10 panel(s) C provides comparisons of SLA variance reduction including the standard VBayes and OLS when used to correct data from pass 003 of the Cal/Val orbit in and around the United Kingdom. It can be seen that the OLS provides the greatest reduction in SLA when considered without uncertainty. This fact is unsurprising, as the estimator is simply finding parameter values which maximally describe the dataset. Inspection of the standard VBayes and Spatial VBayes as the confidence level is increased tells a different story. While considerable tidal variability can be found in the SLA in and around the complex UK coastline, residual tides either do not exist or, more likely, cannot be determined from the available data in regions shown in blue. An example of this is shown in the bottom right panel $(50.222°N 1.022°W)$. Without considering uncertainty, both OLS and VBayes tidal estimation suggest the observed SLA variance can be reduced by around 40%. While some tidal variability may remain, it can also be seen that both estimators are biased by the nonstationary and clearly nontidal SLA spikes which were not filtered out. All values in the original SLA can be seen to fall within the spatial VBayes $2\sigma$ confidence and thus, we cannot conclude at the 95% confidence level that the signal contains any tidal variability. We believe this interpretation provides a simple and statistically grounded approach to estimate the residual tidal variability in the corrected SLA products. Using this approach, Table 3.1 estimates the residual tidal variability for the L2 2 km Cal/Val SLA in passes 003 and 016 between 50 and 60 degrees latitude. Using this approach we estimate with a confidence level of 95%, on average 7.2% of the SLA variance in this region is attributable to tidal variability. The slight increase in estimated tidal variance between Standard VBayes and Spatial VBayes is due to the fact that spatial coherence yields a reduction in epistemic uncertainty (e.g. uncertainty which arises from insufficient information). This makes sense conceptually as when the information contained in adjacent points is congruous with the central point (spatially coherent), our model has greater confidence in its prediction. The 7.2% estimate should be viewed as a conservative lower bound as we are only fitting 8 harmonics and are largely neglecting diurnal tides. Regardless, we believe our approach when combined with the high-resolution KaRIn data is useful in estimating the spatial deficiencies in global tidal models.

For researchers employing the corrections directly, the final choice of confidence level can be selected depending on the nature of their usage. It is expected that additional data from the SWOT science phase will reduce the epistemic uncertainty and help estimate some of the smaller residual tidal signals. Errors in SWOT data processing can yield artifacts in both the estimated tides and the SLA variance reduction estimates. In particular, we see that large SLA variance reductions can be seen in and around the nadir gap. Reviewer comments have suggested this is a consequence of incorrect remapping from the SWOT measurement grid to the fixed reference grid in the data products. Irrespective, we find that the pruning procedure provided seems to effectively reduce these errors in the corrected SLA.

## 3.5 Discussion

Arguably as important as correcting errors from present geophysical corrections is estimating where and how they are deficient. Despite their remarkable performance and reliability, many global tide models, FES included, do not provide uncertainty estimates (19). The VBayes residual analysis of SWOT data provides a potential solution to this limitation by approximating these uncertainties directly, including those created by spatial interpolation. Critically, this provides a unique high-resolution view into both the short-scale spatial characteristics of tides and, by extension, the errors made by tide models. As demonstrated, equivalent estimates cannot accurately be obtained using conventional OLS harmonic analysis.

One important feature of our method is the novel incorporation of spatial coherence outlined in Section 3.2.4. This approach formalises a long standing practice of binning altimetric measurements for empirical tidal models (87). While our approach also weights measurements according to their distance to the central node we provide several further developments. First, we add an additional weighting based on the probability that the observations at the given point are correlated with those of the central point. This is particularly important in complex coastlines where greater spatial variability is present. Second, rather than assume the quadrature components of the central node describe all neighbors within the given radius, we instead allow a linear variation between these nodes. This new approach shows significant improvement for simulated and independently derived SWOT data. The inductive bias effectively increases the ratio of measurements to parameters by a factor of two when using just four neighbors and can scale with additional points. We have shown that this procedure can help reduce aliasing errors and improve the resolution of minor tides. Future work is needed to determine the breakdown point of the linearity assumption when the radius of neighboring points is increased beyond 4 km. Based on this, the spatial coherence can be applied to ungridded data from historical altimetry missions.

We wish to stress that no model is a silver bullet, and the Spatial VBayes is not without its limitations. In low noise regimes, the OLS is by very definition, a superior estimator. Furthermore, assumptions made by the variational Bayesian approach and the epistemic uncertainty from the limited measurements can lead to underestimation of the true uncertainty. The construction of our methodology has made every attempt to mitigate these limitations. We advise that uncertainty estimates should be viewed as a lower bound on the true parameter uncertainty. As such, the final choice of confidence level is left to the end user to allow researchers to decide how confident they wish to be depending on their application. It is also worth remarking that for a majority of regions, assimilative tidal models such as FES2014b perform phenomenally well (23). In these regions empirical correction of the SLA can actually introduce error into the signal. The procedure outlined in Section 3.4.2 helps to mitigate this by allowing a user defined threshold of residual tidal variability in order to apply corrections. For example, by simply requiring that a location's SLA contain at least 5% tidal variance at a 95% confidence level, testing with the gauges showed many of these errors can be avoided.

Notably, the VBayes method does not explicitly consider temporal correlations in the residual, though it provides increased robustness to these errors compared to OLS. Due to the background red-noise intrinsic to the noise spectrum this can lead to inaccurate parameter estimation. While several approaches were considered to account for this variability such as considering a red-noise covariance (67), the sparsity of the SWOT data significantly degraded the effectiveness of approaches which rely on spectral estimation of the residual. This represents an important area of future work when applying VBayes to more conventionally sampled reference series. In this vein, the VBayes estimator can be modified in several ways. Perhaps the easiest is to treat the residual as a mixture of Gaussians, which can overcome the assumption of homoskedastic residuals (68). Additionally, to our knowledge, regularization has not yet been formally considered within the context of frequency resolution in tidal analysis. While several authors have experimented with simple Ridge regression models (88; 89), theoretical work is needed to formalize these improvements relative to the standard Rayleigh and Munk-Hasselman criterions and how it relates to the characteristic failure mode of unregularized least-squares as record length is decreased.

Future work will look to apply the Spatial VBayes approach to the more than 30 years of historical altimetry measurements now available. As demonstrated in both the simulated and real validation, the reduced sensitivity to noise, outlier removal, and uncertainty estimation are useful properties within the context of SLA residual analysis. As the science phase of the SWOT mission nears completion of its first year, sufficient data will exist to produce analogous analyses at a global scale. With the approximately 21 day repeat orbit of this phase, the advantages garnered by the Spatial VBayes will prove even more salient. Additionally, the VBayes estimator has shown very interesting properties regarding tidal super-resolution (e.g. resolving constituents which are not resolvable by the conventional Rayleigh criterion) and overcoming “aliasing”, as illustrated by the example in Figure 3.1, and the improved resolution of other constituents for the sun synchronous Envisat orbit in Appendix A. To be explicit, no-degree of super-resolution can overcome the aliasing of the S2 constituent for sun-synchronous missions. However, as is shown, unregularized estimators (OLS) can exhibit bias across other parameters if the aliased parameter is included in the analysis. Hence, the ARD priors of the VBayes help prevent the errors from S2 estimation from degrading the estimation of other constituents. Future work will look to formalize this result as it could be useful in maximizing the scientific return from historical sun synchronous missions.

Finally, it remains to be said how the SWOT data themselves can actually be used to further our scientific understanding of tides. Early analysis of the high-resolution data products in the Bristol Channel has revealed remarkable spatial characteristics of compound and nonlinear tides. This presents the opportunity to study features such as tidal asymmetry, and tidal interactions with river outflow. The unique 1-day repeat orbit of the SWOT Cal/Val orbit also allows for the study of relatively short-scale temporal processes. Future work will look to study the temporal and spatial characteristics of storm surges using these data. These applications are only scratching the surface of what is possible using wide-swath altimetry. However, we note that the success of these applications remains dependent on the quality of the tidal corrections themselves. It is our hope that the Spatial VBayes approach can help to enable this success.

## 3.6 Conclusion

This work develops a spatially coherent harmonic analysis using variational Bayesian inference rather than conventional least-squares estimation. The method demonstrates superior robustness to several types of noise contamination common in altimetric tidal analysis. Several, techniques are developed using the Bayesian uncertainty estimation to automatically identify and remove erroneous measurements. The method demonstrates excellent performance relative to state-of-the-art models such as FES2014b when validated on tide gauges on the European shelf. The largest improvements can be found in the the Elbe Estuary, GER. Further analysis of SLA residual variance reduction confirms these results and provides a first look at swath based uncertainty estimation for FES2014b using the SWOT Cal/Val data. We estimate a lower bound on the residual tidal variability in the corrected SLA with a confidence level of 95% of at least 7.2% on average for Cal/Val passes 003 and 016. Additionally, estimates of the MSS error highlight large uncertainties around complex coastlines arising from tidal flats, sediment deposits, and propagated errors from tidal corrections exceeding 1 meter in regions such as the Bristol Channel, UK.

SWOT Cal/Val Tidal estimates from the Spatial VBayes methodology can serve both as a source of corrections for early SWOT data, and as a unique and independent estimate of the spatial characteristics of the uncertainty in present geophysical corrections. An implementation of the variational Bayesian tidal analysis is provided as an open-source Python package (VTide) and can be employed for conventional tidal analysis and prediction (https://github.com/thomasmonahan/V

## Open Research Section

All low-rate sea-surface height SWOT data is made publicly accessible via `https://doi.org/10.24400/527896/A01-2023.015`. The data used for validation in Sections 3.4.1 and 3.4.2 make use of the pre-validated L2_LR_SSH version 2.0 ( version C). All results for the Bristol Channel, UK make use of the pre-validated L2_LR_SSH version 1.0 ( version B) products. The FES2014b atlas used for tide-gauge validation can be accessed through Aviso (22). We note that we utilized the 'extrapolated' version which yields superior performance in coastal regions. The TICON-3 dataset used for synthetic validation is described in (52) and can be accessed at: `https://doi.org/10.1594/PANGAEA.951610`. All code required to reproduce each analysis is available at `https://doi.org/10.5281/zenodo.14892032`. The code is also maintained as the VTide Python package: `github.com/thomasmonahan/vtide`

## Metrics

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

1.

$$\mathrm{RMS}_k = \sqrt{[(A_{\mathrm{model}}\sin(\omega_k t) + B_{\mathrm{model}}\cos(\omega_k t)) - (A_{\mathrm{true}}\sin(\omega_k t) + B_{\mathrm{true}}\cos(\omega_k t))]^2} \tag{3.23}$$

2.

$$\mathrm{RSS} = \sqrt{\sum_{k=1}^{n} (\mathrm{RMS}_k)^2} \tag{3.24}$$

3.

$$Var_{diff} = 100 \times \left(\frac{Var_{FES} - Var_{model} - N\sigma_{model}}{Var_{FES}}\right) \tag{3.25}$$

4.

$$\text{F1-Score} = \frac{\mathrm{TP}}{\mathrm{TP} + 0.5(\mathrm{FP} + \mathrm{FN})} \tag{3.26}$$

5.

$$\mathrm{SNR}\ (\mathrm{dB}) = 10 \cdot \log_{10}\left(\frac{\sigma^2_{\mathrm{signal}}}{\sigma^2_{\mathrm{noise}}}\right) \tag{3.27}$$

# 4 Observations of the seiche that shook the world

# Abstract

On September 16th, 2023, an anomalous 10.88 mHz seismic signal was observed globally, persisting for 9 days. One month later an identical signal appeared, lasting for another week. Several studies have theorized that these signals were produced by seiches which formed after two landslide-generated mega-tsunamis in an East Greenland fjord. This theory is supported by seismic inversions, and analytical and numerical modeling, but no direct observations have been made. Here, we present primary observations of this phenomenon using data from the Surface Water Ocean Topography mission. By ruling out other oceanographic processes, we validate the seiche theory of previous authors and independently estimate its initial amplitude at 7.9 m using Bayesian machine learning and seismic data. This study demonstrates the value of satellite altimetry for studying fast oceanic processes and extreme events, while also highlighting the need for specialized methods to address the altimetric data's limitations, namely temporal sparsity. These data and approaches will help in understanding future unseen extremes driven by climate change.

## 4.1 Introduction

Extreme events are evolving as a direct consequence of climate change, leading to the emergence of new, previously unobserved phenomena (90; 91). In remote regions like the Arctic, where in-situ measurements are sparse, scientists depend on analytical and numerical models to explore these events. However, modeling in such regions presents significant challenges due to the uncertainties in the data required to calibrate and validate these models (92). Consequently, large simplifications are often necessary, potentially resulting in substantial discrepancies between observed and modeled phenomena.

The mysterious 10.88 mHz very-long-period (VLP) seismic signal, which appeared following a tsunamigenic landslide in the Dickson fjord, Greenland, on September 16th, 2023, and the subsequent interdisciplinary scientific efforts to determine its origin, underscore these challenges. Two independent studies (93; 94) have hypothesized that the signal was driven by a standing wave, or seiche, which formed in the aftermath of the tsunami. While it is well-documented that seiches can form in resonant enclosed and semi-enclosed basins (95), the loading-induced tilt they produce has only been observed locally ($< 30$ km) and for short durations ($< 1$ hour) (94; 96). Moreover, no prior evidence exists of persistent fluid sloshing (lasting several days) without an external driver. The 9-day attenuation of the globally detected VLP seismic signal is thus highly anomalous. Even more curious was the reappearance of the signal on October 11th, 2023, as noted by (94), with approximately half the magnitude and duration of the initial event. This recurrence coincided with a second tsunamigenic landslide in the same gully in the Dickson fjord.

The rock-ice avalanches and subsequent tsunamis from both events have been well-documented through a combination of satellite and field observations, including evidence of tsunami run-up within the fjord and as far away as the research station at Ella Ø, over 72 km away (93; 94) (Section 4.4.2.9, and Supplementary Figure 1). In contrast, evidence for the seiche has, by necessity, relied solely on a combination of analytical and numerical models, supplemented by seismic observations. While the seismic amplitudes are well reproduced, significant discrepancies remain between these studies in the estimated initial amplitude of the seiche (ranging from 2.6 m to 7.4–8.8 m). Even more perplexing is the fact that the earth-shaking seiche was not observed by the Danish military whilst surveying the fjord on September 19th (94). Hence, empirical observations are essential not only to confirm the existence of the seiche but also to validate the models and refine our understanding of the event dynamics. To support the conclusions that the 10.88 mHz VLP signal was produced by a seiche in the Dickson fjord the authors of (93) and (94) applied

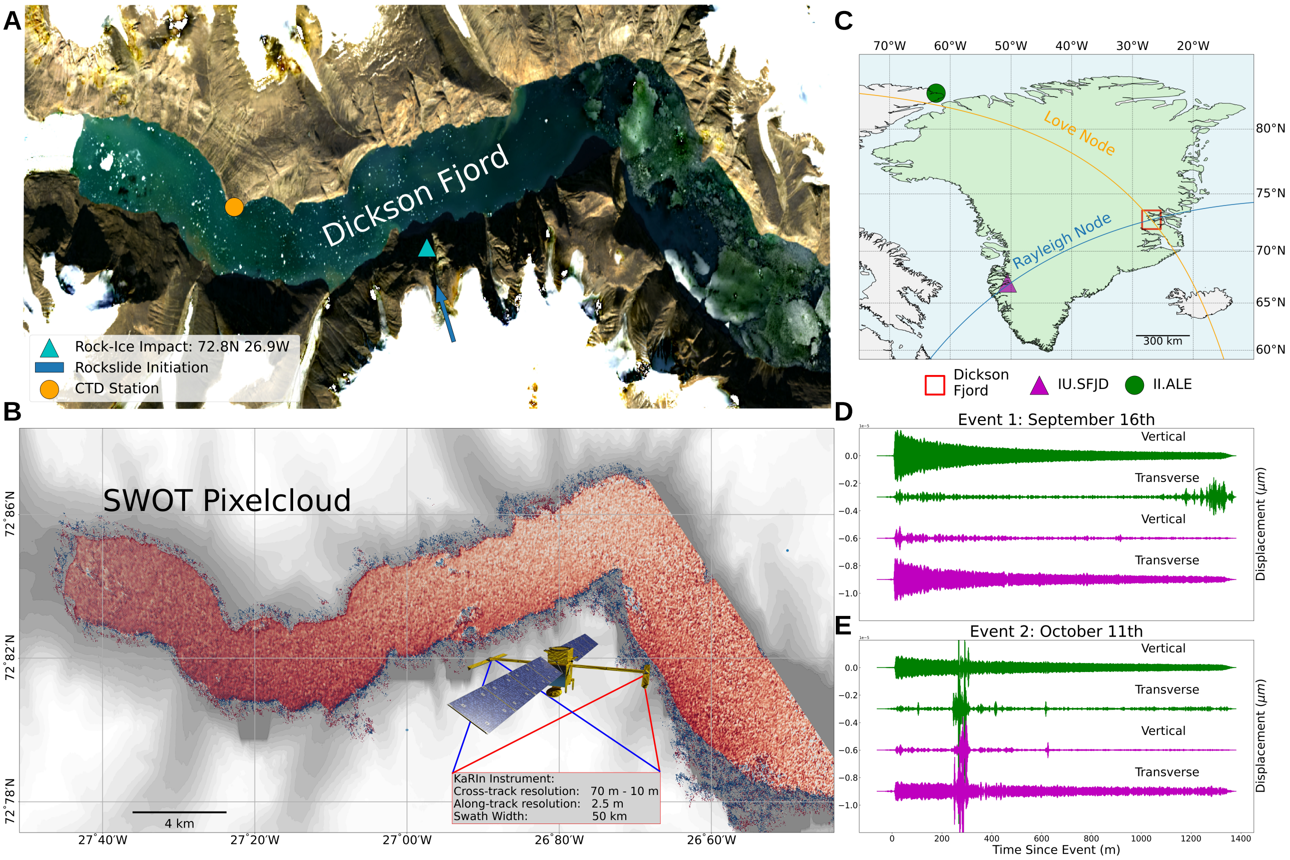

Figure 4.1: Dickson fjord study region, SWOT measurements, and in-situ measurements. (**A**) Sentinel-2 image of the Dickson fjord in summertime, with rock-slide, CTD tide-gauge, and atmospheric station shown (Copernicus Sentinel data 2023). (**B**) SWOT pixelcloud measurements from a single pass over Dickson fjord, measurements colored by measured sea-surface height. SWOT satellite courtesy NASA/JPL-Caltech. (**C**) Visualization of the study region (Made with Cartopy Natural Earth (2)), and nearby IU.SFJD and II.ALE seismic stations. Rayleigh and Love nodes are plotted in blue and orange. (**D & E**) Seismic observations at II.ALE (green) (3) and IU.SFJD (magenta) (4) bandpass filtered between 10-13mHz for the September 16th (D) and October 11th (E) events. Time units are minutes.

two independent approaches. Both studies identify the observed radiation pattern of seismic Rayleigh and Love waves to be consistent with an oscillating single force perpendicular to the Dickson fjord. Seismic Rayleigh waves exhibit elliptical, rolling motion with both vertical and radial displacement, propagating radially from the source. Love waves exhibit purely transverse motion perpendicular to their propagation direction. Additionally, both groups identify seismic source locations near the Dickson fjord. Here we present a much abbreviated summary of the additional methods and evidence.

By performing a seismic inversion on three sets of teleseismic arrays, the authors of (93) isolate a predominantly horizontal and perpendicular force to the Dickson fjord. This finding serves as the basis of a simple analytical model of the sloshing physics by considering a simplified rectangular fjord geometry of 2 x 20 km. Using the single force inversion, they identify an initial horizontal force of approximately 160 GN which leads to an estimated initial amplitude of 2.6 m. While a 2-d finite difference modelling approach was applied in the simplified rectangular geometry, this model only serves to illustrate that the fundamental-mode oscillation can form. Much effort is devoted to identifying the physical drivers of the observed damping, which, due to our approach not requiring these, we do not discuss further.

The large interdisciplinary team in (94) use a combination of high-resolution numerical simulation and analytical modelling to corroborate their claims. Two numerical approaches are considered; however, their preferred approach employs a nonlinear hydrostatic model implemented in HySea (97) which treats the landslide as a granular flow. Using a fine grid spacing of 3 m, the simulation stabilizes into a slowly decaying seiche after approximately 5 minutes with an initial amplitude of 7.4 m. Notably, the first eigenmode has an oscillation frequency of 11.45 mHz (85 seconds), which differs from the observed VLP signal frequency of 10.88 mHz (92 seconds). This numerical simulation is then employed as a source time function to generate global seismic waveforms. Through direct comparison of the simulated envelope, the authors find good agreement between the synthetic and observed signal attenuation. An analytical model is also utilized, using a more realistic simplified geometry than in (93). The authors identify the initial force to be 500 GN, significantly larger than the 160 GN estimate in (93). This yields an estimated initial seiche amplitude of 8.8 meters.

While both studies provide compelling evidence that the source of the persistent 10.88 mHz signal was a seiche originating in the Dickson fjord, significantly different values are obtained for the initial amplitude of the seiche (2.6 m vs. 7.4-8.8 m). Both studies attribute these discrepancies to unmodelled effects. We note that both studies consider significantly different simplified fjord geometries e.g. (93) assume a fjord width of 2 km, and a length of 20 km, and (94) assume a width of 2.7 - 2.88 km (depending on the figure) and a length of 10 km. Naturally, these choices will lead to different analytical estimates of seiche amplitudes.

Here, we offer a completely different approach – using primary observations of this phenomenon to answer these questions. Using these data in conjunction with seismic observations, we estimate initial amplitudes of the September and October seiche events and compare these results with the prior studies. These observations and estimates provide conclusive evidence for the existence of this phenomena.

## 4.2 Results

### 4.2.1 Surface Water Ocean Topography (SWOT) Mission

In contrast to in-situ devices, satellite altimetry provides near-global measurements, albeit with an inevitable trade-off in temporal sampling (98). After more than 30 years, these data have revolutionized our understanding of many oceanic and climatic processes (46). However, significant challenges arise in the study of extreme events due to a combination of the temporal sparsity, and the 1-d nature of the observations. Conventional Nadir altimeters sample data directly beneath the spacecraft, producing 1-d profiles along the sea-surface. This sampling severely limits the ability to draw conclusions regarding the spatial dynamics of extremes, and often leaves events unobserved altogether.

The new wide-swath Surface Water Ocean Topography mission, launched on December 15th, 2022, has overcome many of these deficiencies (48). Unlike conventional Nadir altimeters, the KaRIn instrument onboard SWOT provides ultra-high resolution 2-d measurements of ocean surfaces extending 50 km on either side of the spacecraft (99). SWOT provides high accuracy measurements directly up-to coastlines, and uniquely into fjords, with an effective pixelcloud resolution of 2.5 m along-track and a variable resolution ranging from 10 m to 70 m in the cross-track direction (48). These measurements have low instrumental noise of less than 0.4 cm (100). An overview of these sampling characteristics and an example of a single SWOT pass over the Dickson fjord is shown in Figure 4.1, Panel C. The pixelcloud contains more than 300,000 measurements and provides complete coverage of the study region.

After transitioning to the Science orbit phase of the SWOT mission, SWOT made several observations of the Dickson fjord shortly after the occurrence of both tsunamigenic landslides. For the September 16th event, these passes occurred 0.5-days, 1.5-days, and 4.8-days after the VLP developed. For the October 11th event, only a single 'usable' pass existed 0.5 days after the VLP began.

### 4.2.2 Empirical Observations

SWOT pixelcloud observations of the Dickson fjord for both the September and October events are shown in Figure 4.2. For the October 12th observation (0.5 days post-tsunami), a large negative cross-channel slope (relative to the line $\overline{X_1X_2}$) can be observed across the minor-axis of the fjord. Here we refer to the longitudinal-axis of the fjord as the major axis with the minor axis sitting perpendicular to it. While some noise artifacts exist around 27°W, the spatial distribution of resonant nodes is in good agreement with the high-resolution tsunami simulation in (94).

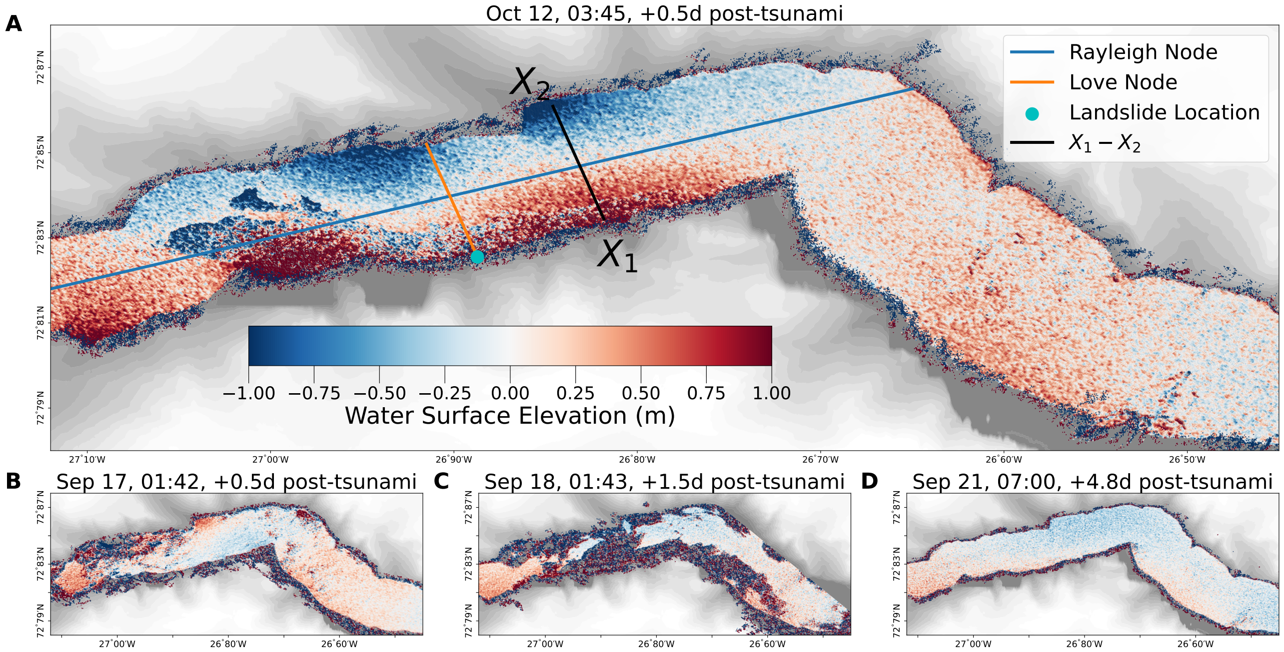

Figure 4.2: Pixelcloud sea-surface elevation maps of the Dickson fjord in the days following the two tsunamis. (**A**) SWOT observation of the fjord 0.5 days after the October 11th tsunami. Rayleigh and Love nodes are overlaid to show the theorized axis of propagation. (**B,C,D**) Consecutive SWOT observations of the fjord 0.5 days, 1.5 days, and 4.8 days after the September 16th event respectively.

Unfortunately, strong noise artifacts muddle large portions of the September 17th and 18th observations. However, on September 17th (0.5 days post-tsunami), a cross-channel slope can be observed antiphase to both the October 11th and September 18th observations. Due to the noise artifacts on September 18th, estimates of the slope are not accurate. A description of the possible sources of these artifacts is given in Section 4.4.1. The September 21st observations exhibit almost no noise artifacts and a very shallow negative cross-channel slope. The spatial distribution of nodes is consistent with those observed on October 12th and the tsunami simulation in (94). We note that significant discrepancies exist between the simulated and observed sea-surface, particularly in the upper reaches of the fjord. Here, the SWOT observations suggest the presence of a persistent water build-up which appears to go unmodeled due to boundary choices.

### 4.2.3 Seismic Attribution

The SWOT data alone cannot estimate the total amplitude of the seiche as the observations could have occurred at any phase of the oscillation period. To overcome

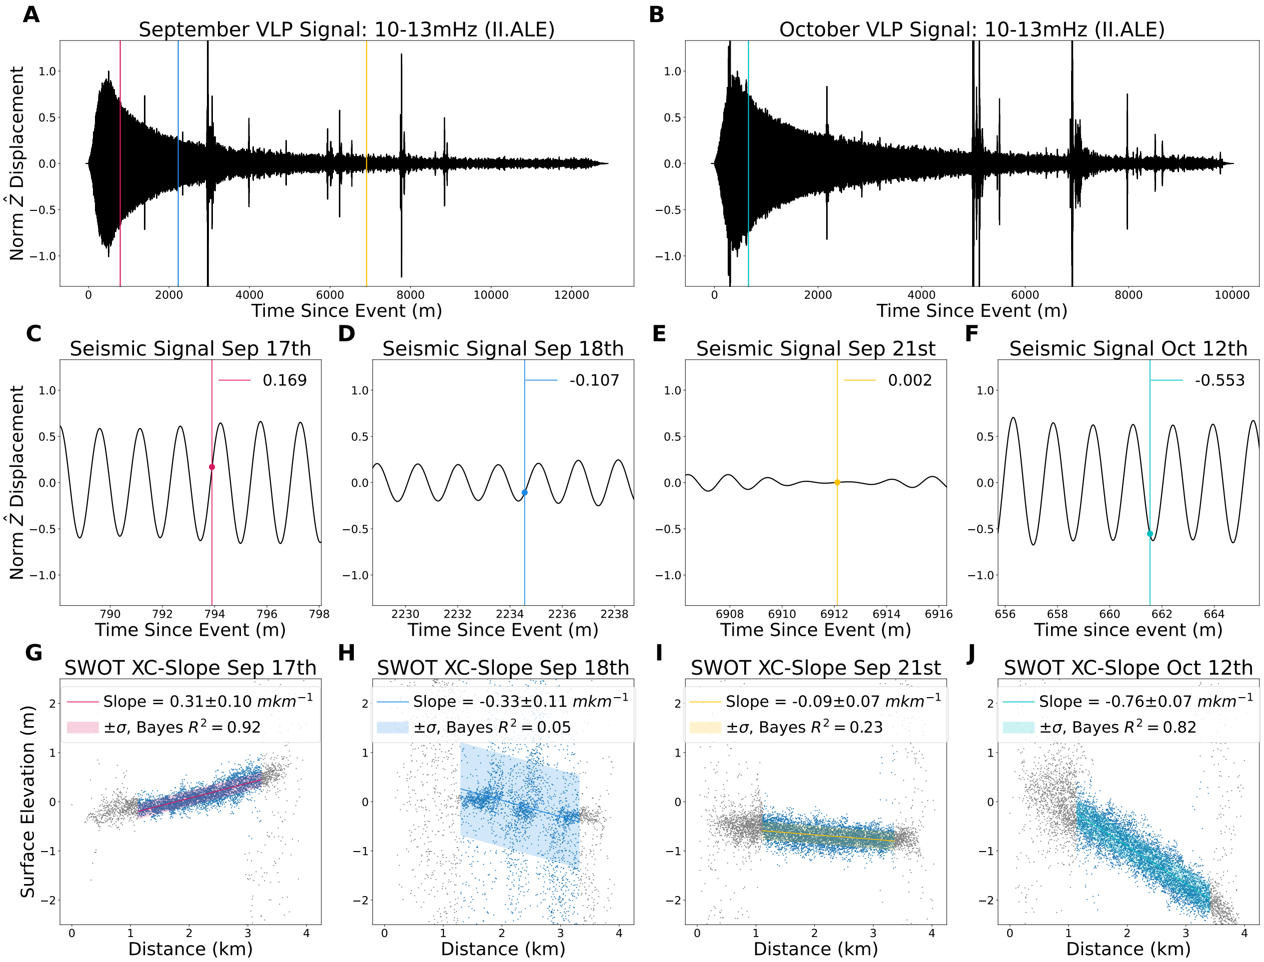

Figure 4.3: Seismic observations of September and October VLP signals at II.ALE Seismic Station and SWOT cross-channel slopes. (**A,B**) Normalized vertical ($\hat{Z}$) displacement VLP signals filtered between 10-13 mHz for the September and October events respectively. SWOT observations are given by vertical lines. (**C-F**) Normalized vertical ($\hat{Z}$) displacement VLP signals with signals observed by SWOT shown as vertical lines. Observed magnitudes relative to the maximum amplitude are shown. (**G-J**) Corresponding SWOT cross-channel observations from $X_1$ to $X_2$. Slope estimates and associated Bayesian $R^2$ values from a Bayesian linear model are provided (Section 4.4.2.2).

this, SWOT observations are referenced to seismic data from the II.ALE seismic station located at 82.5033°N 62.3500°W (1322.9 km away). The II.ALE station sits directly adjacent to the Love node and is thus characterized by almost exclusively Rayleigh waves as shown in Figure 4.1, Panels B, and D. The measured vertical displacement at II.ALE is filtered between 10 and 13 mHz, and shown for both events in Figure 4.3, Panels A and B. SWOT observations are shown as vertical lines. We estimate a phase speed of 4.03 $\mathrm{km\,s^{-1}}$ using a heterogeneous Earth model (101), and an approximate distance from the seismic source of 1409.5 km computed using the same heterogenous Earth model (see Section 4.4.2.1 for details, and Appendix B Figure B.1 for relative location). Estimates of the uncertainty in this value, as well as validation using synthetic data are also provided in Section 4.4.2.1. The relative magnitude and phase of the seiche can be directly determined through comparisons of the SWOT observed cross-channel slopes with the observed ground motion. Snapshots of the observed vertical VLP signal with SWOT observations highlighted are shown in Figure 4.3, Panels C-F. The SWOT observed cross-channel slopes, computed between points $X_1$ and $X_2$ which gave rise to these signals are plotted directly below in Figure 4.3, Panels G-J. The observed cross-channel slopes nicely correspond to the vertical displacement produced at station II.ALE. That is, negative cross-channel slopes (from $X_1$ to $X_2$) are associated with a negative vertical displacement and vice versa. The magnitudes also show good agreement. We note that this is exactly what is expected from the horizontal force produced by the seiche oscillation (Section 4.4.2.6) (94).

To validate these observations, the normalized VLP signal is used to estimate the initial amplitude of the seiche. Estimated slopes are computed using a Bayesian linear model (Section 4.4.2.2). Uncertainty estimates are obtained at both the parameter level and with regard to the noise content of the data. Due to the fact that each study assumes a different width and length of the fjord, we instead consider each estimate in terms of the corresponding cross-channel slope at maximum amplitude (MXCS). For (94) their initial amplitude estimate of 7.4-8.8 m translates to an MXCS of 2.56-3.13 $\mathrm{m\,km^{-1}}$, and the initial amplitude estimate of 2.6 m by (93) yields an MXCS of 1.3 $\mathrm{m\,km^{-1}}$. Using the SWOT observations from September 17th we estimate the MXCS to be $1.83 \pm 0.59$ $\mathrm{m\,km^{-1}}$. This value sits in between the estimates of both previous studies, however, the analytical estimate by (93) sits at the very bottom of the confidence intervals, and the numerical estimate from (94) just outside the top. Data from September 18th and 21st were not utilized due to large relative uncertainties and low Bayesian $R^2$ scores (102) (see Section 4.4.2.2). For the October event, we estimate the MXCS to be $1.37 \pm 0.13$ $\mathrm{m\,km^{-1}}$. The fact that the October 12th observations occurred near a local minimum in the seiche's oscillation allows for tighter uncertainty estimation.

### 4.2.4 Ruling out other suspects

While the SWOT data provides an unprecedented look at the instantaneous water levels in the Dickson fjord, it is only that; a snapshot. The observed cross-channel slopes conform to our expectation of a standing wave oscillating perpendicular to the major-axis of the fjord. However, there are other geophysical phenomena which can give rise to large cross-channel slopes in enclosed basins, namely tides (103) and wind-driven circulation (104). Consideration is given to each of these possible causes.

#### 4.2.4.1 Tides

While tides are obviously not the source of the approximately 92-second seismic signal, due to the sparsity of the SWOT measurements, tidally driven cross-channel slopes could lead to false conclusions about the presence of a seiche in the SWOT data. As there is only one tide gauge in the region, insufficient in-situ information exists to rule this out. Additionally, small fjords are poorly resolved by state-of-the-art global tide models due to interpolation and can thus not be relied upon (28; 23). We instead directly analyze the SWOT pixel cloud data obtained over the year following the two events using a spatially coherent Bayesian harmonic analysis procedure developed in Chapter 3 (105). This procedure has been shown to improve tidal estimation over conventional least-squares approaches for sparse reference se-

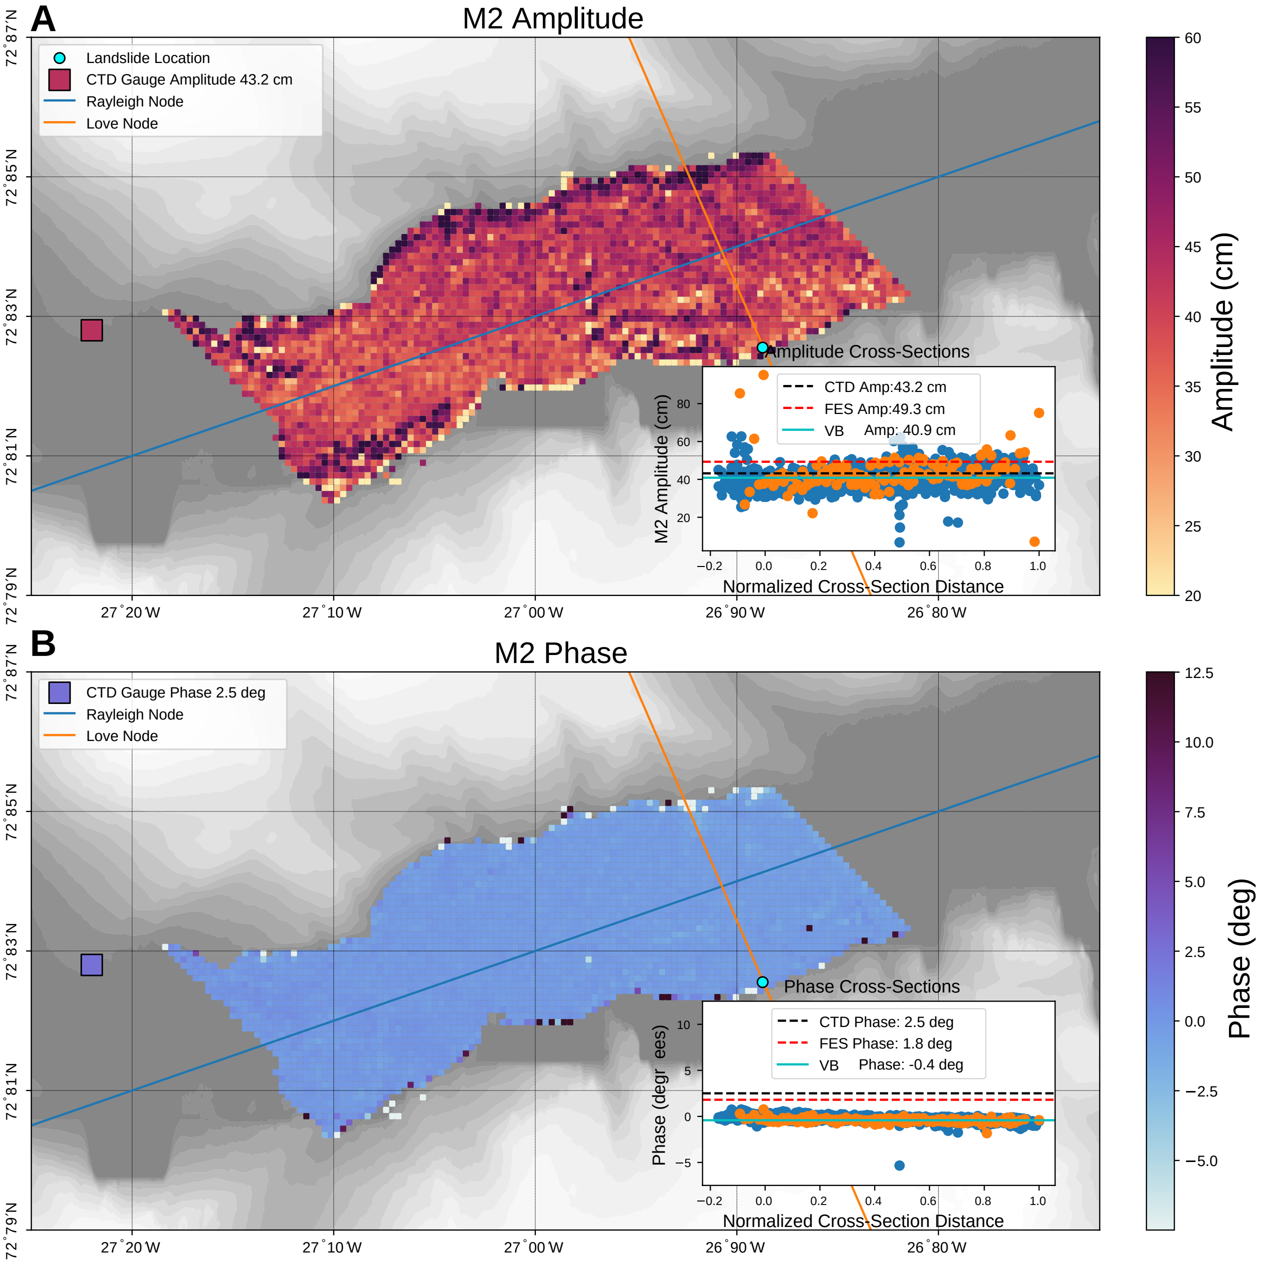

Figure 4.4: Estimates of the dominant lunar tide, M2, from SWOT pixelcloud data using a spatially coherent variational Bayesian harmonic analysis. (**A**) Estimated amplitudes, and (**B**) corresponding phase lags. Estimates are only made for points which have at least 23 measurements. Inset plots show how the amplitude and phase vary along the Rayleigh (blue, left to right) and Love (orange, bottom to top) nodes respectively. The M2 amplitude and phase lag computed from the depth measurements at the CTD station are shown. Horizontal lines show the average VB, FES, and CTD estimates

ries and in complex coastal regions. A complete description of our implementation is given in Section 4.4.2.4.

Figure 4.4 shows empirical estimates of the amplitude and phase lag of the dominant lunar ocean tide, M2. Pixelcloud data is obtained and analyzed from October 20th, 2023 to November 1st, 2024. Due to the presence of winter sea-ice, and errors in the SWOT data, we identified only 23 usable passes. This limited data is insufficient to resolve additional tides; however, it suffices for M2. As expected, both amplitude and phase exhibit small linear trends along the major-axis of the fjord and estimates are in good agreement with the CTD gauge (Figure 4.4). Along the cross-channel direction, a slight rise in amplitude (5-10 cm) is seen in the center of the fjord but tapers off closer to the shores. This behavior may influence the curvature of the cross-channel profile, but will not induce the linear slope observed. The derived phase lags are uniformly distributed along the minor-axis of the fjord. Due to the fact that cross-channel variations in phase lag give rise to cross-channel slopes, we conclude that the observed slope in the SWOT data is not tidally driven.

As noted in (93), an approximately 6 hour modulation of the seismic signal can be observed. While this appears tidally correlated, the behavior is actually opposite of what is expected from a simple traveling wave. A proposed mechanism for this is given in 4.4.2.10.

#### 4.2.4.2 Wind Driven Transport

Wind-driven transport must be considered as a potential driver of the observed cross-channel slopes. For this we must rule out two relevent processes; Ekman transport and wind setup. Consideration to each process is given below.

Ekman transport occurs as a consequence of sustained wind-stress (106). In the Northern Hemisphere, the Coriolis effect causes water to propagate 90° clockwise to the incident wind direction. To assess the possible role of Ekman transport, we consider the Rossby radius of deformation, which defines the horizontal length scale

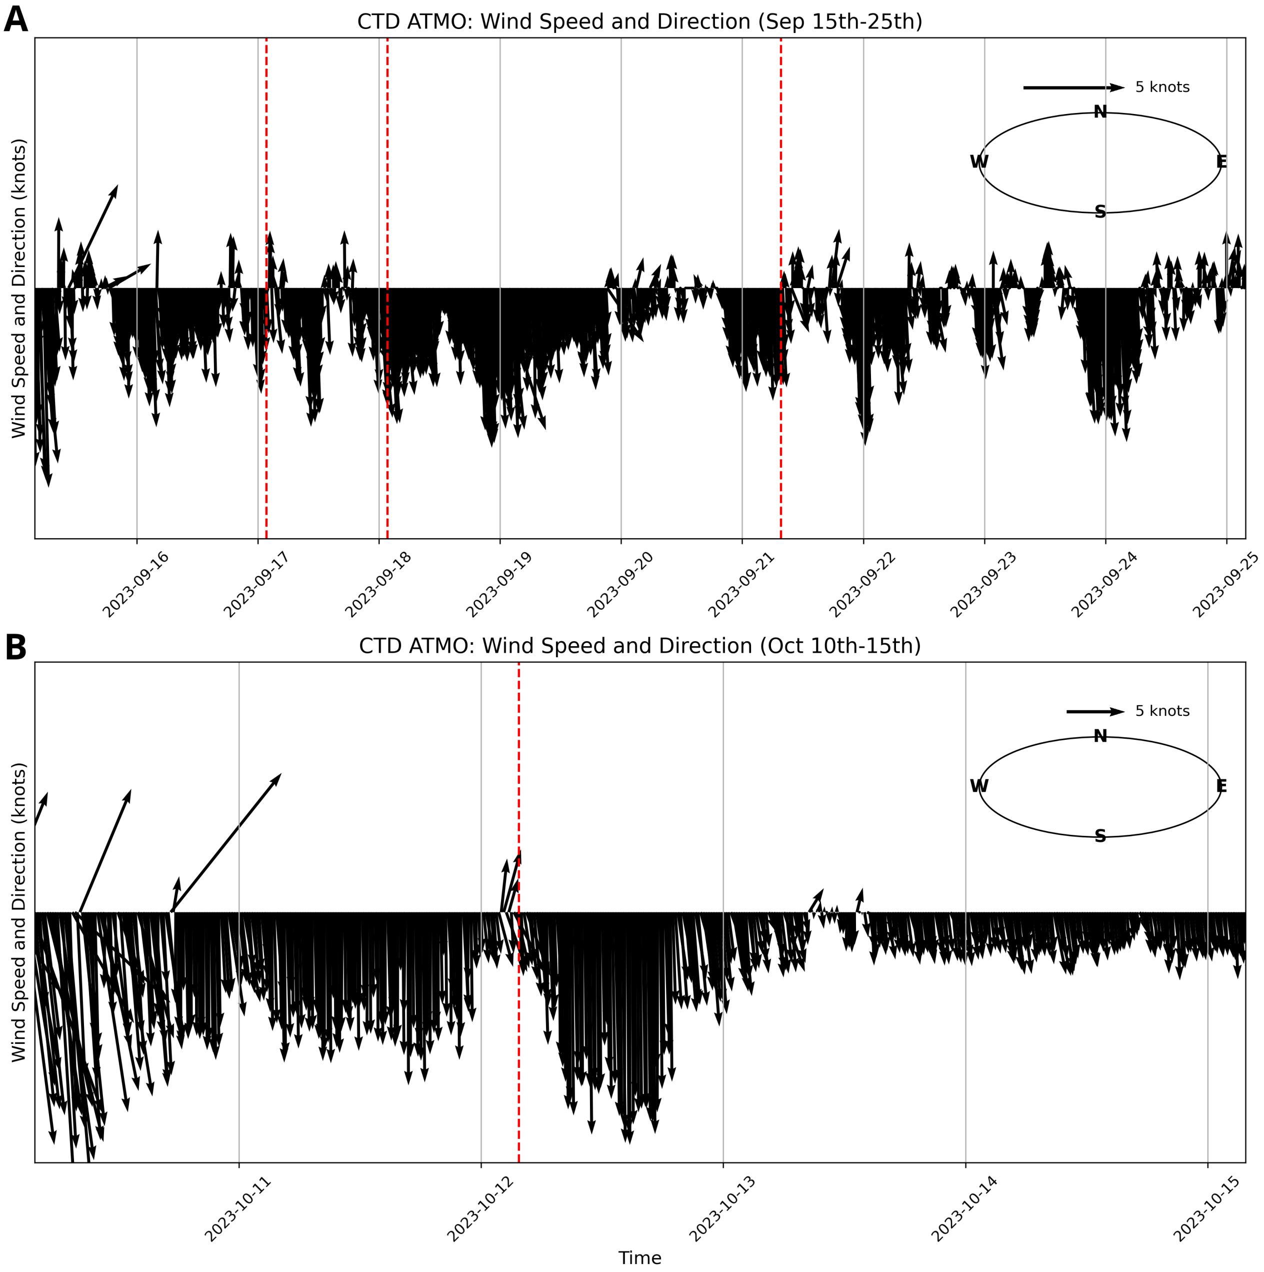

Figure 4.5: Wind speed and direction from the Dickson fjord CTD station during each event (**A**) Wind speed and direction over the duration of the September 16th VLP signal. (**B**) Wind speed and direction over the duration of the October 11th VLP signal. The red dashed vertical lines indicate when SWOT observations occur.

at which rotational effects become relevant in a flow (107). For barotropic flows, the Rossby radius $L_R$ is defined as:

$$L_R = \frac{\sqrt{gD}}{f} \tag{4.1}$$

where $g$ is the gravitational acceleration ($9.81\,\mathrm{m/s^2}$), $D$ is the total depth of the water column, and $f$ is the Coriolis parameter. At the latitude of this study (72.5°N), $f \approx 0.000139\,\mathrm{s}^{-1}$, and for a representative depth of $D = 500\,\mathrm{m}$, we obtain $L_R \approx 500\mathrm{km}$. This barotropic Rossby radius is vastly greater than the ~2.8 km width of the fjord, indicating that rotational effects are not dynamically important at this spatial scale and thus Ekman transport is not a viable mechanism for the observed cross-channel slopes.

For baroclinic flows, which depend on stratification and internal wave dynamics, the Rossby radius is typically smaller. Observations and modeling in Arctic regions suggest that the first baroclinic Rossby radius is approximately 12 km, and increases with stronger stratification (108). However, even this reduced scale is still much larger than the fjord width (2.8 km). We therefore conclude that neither barotropic nor baroclinic Rossby deformation scales are compatible with significant Ekman dynamics in this system.

Wind setup can also be ruled out through evaluation of the wind-speed and direction from the CTD atmospheric station shown in Figure 4.1 Panel A. Results are shown in Figure 4.5. Observations over the duration of the first VLP signal in September show that all SWOT measurements occur after sustained periods of northerly winds at approximately 5 knots. This wind-stress could induce a build-up of water in the Southern portion of the fjord and must be considered closer. However, using a simplified balance from the shallow water equations, we find that a 5-knot wind sustained across the 2.8 km fjord induces a surface displacement of less than 2 mm which sits within instrumental noise levels for SWOT. The observations on October 12th follow yet another sustained period of northerly wind, this time with a higher magnitude of approximately 10 knots. The SWOT observations occur as the wind is switching direction at low magnitudes (< 5 knots). Hence, due to the low magnitude of these winds and theoretical surface displacement of < 2mm, it is unlikely that wind setup contributed significantly to the large (2 m) cross-channel slope observed.

## 4.3 Discussion

Our study provides direct observational evidence of the seiche in the Dickson fjord. Based on the seismic attribution, and systematic ruling out of other dynamic phenomena, we conclude that the observed variability in the SWOT data is consistent with that of a slowly decaying seiche. Thus, this study corroborates the numerical and analytical evidence given in (93) and (94), that the globally observed VLP signal which originated on September 16th, 2023, was due to a seiche which formed after a megatsunami. Additionally, we conclude that the smaller VLP signal observed on October 11th, was also a seiche originating from a smaller tsunami in the very same fjord. Notably, a short period seiche persisting for several days without an external driver has not previously been observed.

Using an empirical and completely independent approach, we estimate the maximum cross-channel slope (MXCS) of the September seiche to be $1.83 \pm 0.59$ m km$^{-1}$. This value falls between the analytical estimate of (93) at 1.3 m km$^{-1}$ and the numerical/analytical estimates of (94), which range from 2.56 to 3.13 m km$^{-1}$. Due to the relatively large uncertainty, our empirical estimate–when considered in isolation–provides limited insight into the true initial value. However, our analysis of the October VLP signal produced an estimate of the October seiche's MXCS at $1.37 \pm 0.13$ m km$^{-1}$. The smaller uncertainty in this case, owing to the October 12th observation occurring near a local minimum of the vertical displacement, gives greater confidence in this estimate.

Seismic data from the II.ALE and IU.SFJD stations show that the initial magnitude of the October VLP signal was approximately twice that of the September event (3 µm vs. 1.5 µm at II.ALE). Since the horizontal force is proportional to the average cross-channel slope (see Section 4.4.2.6), we made a second estimate of the September MXCS to be $2.74 \pm 0.26$ m km$^{-1}$. This value closely aligns with the numerical/analytical estimates of 2.56–3.13 m km$^{-1}$ reported by (94). Therefore, based on the relative magnitudes of the two forces at adjacent seismic stations, and the robust estimate of the October 12th event, we conclude that the numerical and analytical estimates by (94) are in good agreement with the real data. Furthermore, we argue that the estimate provided in (93) likely underestimates the true magnitude due to inaccuracies in the assumed fjord geometry, and the underestimation of the initial force at only 160 GN. If we consider a mean fjord width of 2.88 km as in (94), our empirical estimate suggests the tsunami stabilizes into an initially 7.9 meter tall seiche.

This study highlights the value of wide-swath satellite altimetry in characterizing extreme events and oceanic phenomena more generally. Due to their short-periods, seiches have long been difficult to study from conventional altimeters. The spatial resolution of SWOT provides new opportunities in this area, as well as for studying other fast-moving oceanic processes such as storm surge, and even large waves. As shown, these data also provide the opportunity to connect and understand the complex interactions between climate change and the different components of the geosphere. However, this work also emphasizes the importance of specialized, interdisciplinary methods to address the intrinsic limitations of these data, particularly the challenges posed by temporal sparsity. The SWOT data at its present level of processing is not trivial to work with. Dedicated efforts are needed to improve reprocessing of these data in fjords. Additionally, open-source tools that bridge the gap between raw data and analysis pipelines are essential for enabling non-expert users to utilize these data.

While sufficient observational and bathymetric data existed to recreate the observed seiche dynamics numerically, for many remote regions this is not the case (92). Indeed, while the effects of climate change are felt globally, the largest and fastest changes are often in these regions (109). As such, while we echo the claims from previous authors about the need for more in-situ sea-level gauge sensors, we believe a follow-on to the SWOT mission and investment in future wide-swath and ideally non-sun-synchronous altimetry missions is also critical to monitoring these effects. Above all, designing orbits to achieve higher temporal sampling at lower latitudes should be prioritized as it presently limits the study of extremes. Furthermore, we stress that a key tool for identifying these events is the accurate computation of sea-level anomalies (SLA) by applying different geophysical corrections. Due to the complex and narrow geometries of fjords, and a lack of historical altimetric measurements, tidal estimates in these regions are poor (28). Additionally, we find that even when accurate SLA estimates can be obtained, the mean SLA of the fjord does not exhibit signs of seiching. This underscores the importance of accounting for small-scale spatial variances which SWOT and future wide-swath altimetry missions have the potential to address.

## 4.4 Methods

### 4.4.1 SWOT Data and Processing

Pixelcloud data from the SWOT mission are obtained through open Earth Access (last accessed on February 3rd, 2025). We utilize the version 2.0, HR pixecloud data with shortname “SWOT_L2_HR_PIXC_2.0” in the earthaccess API. The SWOT high-rate pixel cloud data is also available at: `https://doi.org/10.5067/SWOT-PIXC-2.0`. At the time of access, the level 2.0 data constitutes the highest processing level available. The Dickson fjord landslides in September and October of 2023 occurred shortly after the transition of SWOT to the Science phase. The SWOT Science phase is characterized by a 20.86 day repeat orbit with 10-day subcycles. The orbit is at an inclination of 77.6°, and is thus non-sun-synchronous, which reduces tidal aliasing. A consequence of the 10 day sub-cycles is that repeat measurements occur in groups leading to relatively short measurement gaps at some portions of the cycle, and complete gaps in the other. Data for the September and October events were obtained by filtering all available passes between September

16-26, and October 10-18th respectively. We select all passes which fall into the bounding box given in Appendix B Table B.1. For the September period, this yields 5 initial passes occurring 0.5 days, 2 days, 3 days, 4 days and 5 days after the event respectively. Similarly, for the October period, 2 initial passes are identified at 0.5 days and 6 days post-tsunami. We process the SWOT data in two stages, first through manual inspection and then using a standard preprocessing procedure.

**Manual Inspection:** Prior to processing each pass we perform a manual inspection. Due to limitations in the L2 processing algorithms applied to SWOT data, entire passes can be contaminated to the point where the data becomes unusable. An example of this contamination can be seen in Appendix B Figure B1. These errors are visually obvious and are manually flagged and removed through visual inspection. This yields a final set of observations for the September event 0.5 days, 2 days, and 5 days, and only a single observation for the October event taken 0.5 days post-tsunami.

**Data Processing:** All standard geophysical corrections are applied including wet and dry troposphere delays and cross-over corrections. Unlike the Cal/Val SWOT orbit, the cross-over correction for the Science phase is much more accurate due to the significant reduction in time between crossovers. Geoid corrections are applied based on the EGM2008 geoid model. Due to interpolation errors with the default FES2022 tidal corrections, a specialized approach is needed to handle tides. A comprehensive discussion on how tides are dealt with is given in Sections 4.4.2.4 and 4.4.2.5. After applying the geophysical corrections, we filter out all measurement values with sea-surface heights exceeding ±4 m. This threshold was empircally determined following manual inspection of the dataset employed for tidal estimation. Due to retracking issues, spurious observations exist over the fjord which exceed this 4 m threshold and are thus filtered out to avoid biasing estimates. To obtain perceptually uniform sea-surface height maps which can be compared between passes, we set the center of each colormap to be the midpoint of the cross-channel slope.

Due to small inaccuracies with the tidal corrections described in Sections 4.4.2.4 and 4.4.2.5, we found this procedure to be more robust.

**Cross-Channel Slope Computation:** Due to the presence of noise artifacts in some of the SWOT measurements, we rely on estimates of the cross-channel slope between the points $X_1$ and $X_2$ for our analysis. We note that the slice between these points is perpendicular to both the long-axis of the fjord, and the Love node shown in Figure 4.2. As a consequence of the nonuniform sampling of the pixelcloud data, it is necessary to define a tolerance which defines how far a point can be from the defined cross-section to be included. We experimented with several tolerance values but converged on a value of 333 m on either side of the defined cross-section. An example of this is shown in Appendix B Figure B3, with approximately 8000 measurements shown. This choice provides a representative 666 m swath ($\approx 5\%$ of the fjord length) with which to estimate the variability/uncertainty of the cross-channel slope. A description of the Bayesian linear model employed for this procedure is given in Section 4.4.2.2.

**SWOT Data For Tidal Estimation:** Pixelcloud data is also used to estimate the M2 tide throughout the Dickson fjord. Due to the severe aliasing produced by the irregular temporal sampling, we make use of all available data. The presence of winter sea ice creates further difficulties as the inclusion of these data will severely degrade the accuracy of tidal estimation. As before, we locate all passes which intersect with the study region (Table B1) between October 20th 2023, and November 1st, 2024. This yields an initial dataset of 241 passes. After manual filtering and removal of passes containing sea-ice, this leaves only 23 high-quality passes which are employed for tidal estimation. We apply the same geophysical corrections as before. In order to obtain time-series which can be fed to our spatially coherent Bayesian harmonic analysis we bin the nonuniformly sampled pixelcloud data into a fixed grid with a 100 m x 100 m resolution. Experiments were conducted at variable resolutions and this was found to yield a good balance between high resolution, and reduced sensitivity to noise. Data points in each bin are averaged. The total number of data per bin is variable, thus, we restrict our analysis to bins which contain at least 23 measurements. This threshold is empirically defined based on the fact that the usage of fewer than 23 measurements in testing produced spurious tidal estimates in some regions.

### 4.4.2 Seismic Data

Seismic data are accessed using the Federation of Digital Seismic Networks (FDSN) web service client for Obspy (110). As both (93) and (94) provide a comprehensive analysis of the long-period seismic signal, we choose to focus on two representative stations: II.ALE (3), and IU.SFJD (4). In order to isolate the energy associated with the VLP, data are shown after being bandpassed between 10 and 12 mHz. All additional filtering parameters needed to replicate our results are provided in the "seismic_attribution.ipynb" notebook.

#### 4.4.2.1 Seismic Attribution

Seismic attribution is performed using data from the II.ALE seismic station. The station is situated 1322.9 km away from the landslide source at 82.5033°N 62.3500°W. The station sits directly adjacent to the Love node (160°) which runs perpendicular to the major-axis of the fjord (Figure 4.1 Panel B). As a consequence, this station receives almost exclusively Rayleigh waves which is reflected by the dominant vertical component of the VLP signal shown in Figure 4.1 Panel D. Referencing the SWOT observations to these waveforms requires the accurate identification of the phase velocity and source location of the VLP signal (111).

**Phase Velocity:** To compute the Rayleigh wave phase velocity we utilize a heterogeneous Earth model as in (94). To obtain the phase velocity of the 10.88 mHz signal we interpolate from the 10 mHz and 15 mHz LITHO1.0 velocity models (101). Integrating over the path between the Dickson fjord and the II.ALE station we obtain a mean phase velocity of 4.0393 $\mathrm{km\,s^{-1}}$.

**Source Location:** Due to uncertainties in the computed phase velocity, the source location of the VLP signal is not necessarily given by the center of the Dickson fjord. Indeed, both (93) and (94) identify source locations which are nearby, but away from the fjord itself. Using the same heterogeneous Earth model and a Fast Marching method for beamforming, a source location 92.4 km from the landslide location at 72.2°N 25.1°W is identified (94). This source location is approximately 1409.5 km away from the II.ALE seismic station. Combined with the computed phase velocity, this yields a travel time of 348.97 s.

**Validation With Synthetic Observations:** To test our hypothesis that the phase velocity and source location obtained using the heterogeneous Earth model are appropriate for seismic attribution we compare our approach to an independent forward modeling exercise carried out in (94). Using their 3 m HySEA simulation of the seiche as a source time function, synthetic 3-component Green's functions are computed using the Syngine web service and convolved to approximate the displacement signals at the II.ALE station. In order to align the two signals, the authors identify an approximately 350 s shift empirically. This value is in agreement with the 348.97 s travel-time computed using the heterogeneous Earth model and beam forming. This agreement between two completely independent methods confirms the validity of the 348.97 s Rayleigh wave travel time to II.ALE for seismic attribution.

**Seismic Uncertainties:** While our estimated Rayleigh wave travel time is in good agreement with the empirical estimate by (94), this represents an important source of uncertainty in the estimation of the seiche amplitude. No comprehensive approach exists for the quantification of Rayleigh wave phase velocity or beamforming location uncertainties (112). As such, we provide a lower-bound on our uncertainty estimate using the discrepancy between our estimate of 348.97 s and the empirical estimate of 350 s by (94). If we assume a 92-second period of oscillation, this yields an approximate error of 2.12 %. Further uncertainties exist due to the fact that near the point source, the phase is governed by a Bessel function rather than a complex exponential as utilized. Whilst the uncertainty this introduces is difficult to estimate, it is important that this is acknowledged (113). Additionally, testing of homogeneous Earth models yielded different estimates of both source location and velocities, yielding variable results. We found the agreement between these models and the empirically estimated wave travel time to generally be worse and thus did not consider them further.

#### 4.4.2.2 Bayesian Regression

Both cross-channel slope and tidal estimation are performed using a Bayesian linear model. Our selection of a Bayesian approach reflects our objective to accurately quantify the uncertainty in the estimated parameters and the SWOT data themselves. Unlike conventional least-squares estimation and other frequentist variants which provide point estimates of the parameters, a Bayesian approach considers (and computes with) the probability distributions of the parameters. By representing our parameters of interest as probability distributions, the uncertainty associated with parameters in our model is explicitly represented thus informing a calibrated measure of uncertainty over the target variables. As will be shown, the selection of appropriate priors can yield further advantages such as natural parameter shrinkage and increased robustness to noise. Shrinkage refers to reducing the magnitude of parameters which do not contribute to the solution. This favors simpler models unless the data justify additional complexity.

Consider the linear model given by $y_i = w^{\intercal} x_i + \epsilon_i$. Here, $y_i$ is the $i^{th}$ SWOT sea-surface height measurement, $w$ is a vector of the estimated weights (corresponding to the entries of the design matrix $X$), $x_i$ is the $i^{th}$ row of an $M \times N$ design matrix $X$, and $\epsilon_i$ is the residual. Here, $M$ is the total number of measurements, and $N$ is the number of input functions (described below). Linear models of this form are employed for both cross-channel slope and tidal estimation. As will be described, in our Bayesian formulation, the inferred elements of the weight vector $w$ govern the shape of the various distributions used in the model. The parametric form for these distributions, and their justification, is given at length below alongside the variational inference procedure. The tidal estimation procedure is provided in Section 4.4.2.4.

**Cross Channel Slope Estimation:** To estimate the cross-channel slope we perform a standard Bayesian linear regression on the instantaneous sea-surface height cross-sections from each pass. The $i^{th}$ row of the design matrix $X$ is simply given by $x_i = [1, x_i]$, where the entry 1 corresponds to the bias (intercept), and $x_i$ the distance along the line $\overline{X_1 X_2}$. Here, the observation $y_i$ is the corresponding SWOT measurement with distance $x_i$ along $\overline{X_1 X_2}$.

**Bayesian Analysis:** Here we provide a brief overview of the Bayesian model utilized and the variational inference method. A complete derivation of variational inference can be found in (72) and a more detailed exposition of the Bayesian linear model, used here, can be found in the appendix of (71). The basis of our analysis is Bayes' theorem, which provides a framework for updating our prior beliefs $p(\theta)$ about the distribution of all parameters (and hyperparameters), based on a set of observations $Y = [y_0, y_1, \dots, y_N]^T$ and the design matrix $X$. A hyper-parameter is itself a parameter, that governs the probability distribution of another parameter. For example, the mean and variance of a Gaussian are the hyperparameters which define the distribution over another variable. By extension, a hyper-hyperparameter describes a variable that controls the distribution over a hyper-parameter. This yields a posterior distribution which describes our final beliefs over $\theta$ and is given by

$$p(\theta|Y) = \frac{p(\theta)p(Y|X,\theta)}{p(Y)}, \tag{4.2}$$

where $\theta$ is the set of parameters and hyperparameters of the model, $p(\theta)$ is the prior, $p(Y|X,\theta)$ the (data) likelihood, and $p(Y)$ the marginal likelihood which acts as a normalizing term in the inference (as it does not depend upon $\theta$). Here we provide a brief overview of the different components of Bayesian analysis.

**Likelihood:** The likelihood $p(Y|X,\theta)$ describes the likelihood that the observed data occurred, given our model. Here, the model is defined by both the design matrix, and our prior assumptions about the parameters $\theta$. We take the likelihood term to be of Gaussian form, equivalent to a least-squares error assumption. In our analysis, we weight the squared residual between observed $y_i$ and model prediction, by a scalar hyper-parameter $\beta$, which represents the precision, or inverse (co)variance of the noise. Discussion and validation of this modeling decision in the context of SWOT based tidal analysis is provided in (105). This has the effect of weighting how tightly our model should fit the data based on how noisy the data is. This assumption yields a Gaussian likelihood term of the form:

$$p(Y|X,\theta) = p(Y|X,w,\beta) = \left(\frac{\beta}{2\pi}\right)^{N/2} \exp\left\{-\frac{\beta}{2}E_Y(w)\right\} \tag{4.3}$$

where we see the likelihood only depends upon $w$ and $\beta$. The error functional $E_y$ is given by

$$E_Y = \sum_{i=0}^{N}(y_i - w^\intercal x_i)^2. \tag{4.4}$$

**Priors:** Central to Bayesian inference is the usage of priors. Priors are distributions over the parameters included in a model, and reflect our initial expectation of the functional forms and values the parameters should have. Here we describe the choice of these priors for both parameters and hyperparameters, and describe how they impact the resultant model. Under the mean-field approximation, we assume the prior over all parameters $\theta$ can be factorized as

$$p(\theta) = p(w|\alpha)p(\alpha)p(\beta), \tag{4.5}$$

where $\alpha = \{\alpha_j\}$ is a set of hyperparameters which governs the scale of the multivariate Gaussian over the weights $w$. We now treat each term individually.

The model weights $p(w|\alpha)$ come from a zero-mean Gaussian prior, with precisions (inverse variance) $\alpha$. This choice serves two purposes. First, a Gaussian is the least informative distribution for a quantity which can be either positive or negative, and is thus unbiased in this way (71). This is important as both cross-channel slopes and quadrature tidal amplitudes can be positive or negative. Second, weights will only be significantly non-zero if the data requires it. Conventionally this is referred to as an Automatic Relevance Determination (ARD) prior as it induces shrinkage over the model weights which do not significantly aid the model in fitting the data. Using this, for an individual weight $w_j$, the prior has the form

$$p(w_j|\alpha_j) = \left(\frac{\alpha_j}{2\pi}\right)^{1/2} \exp\left\{-\frac{\alpha_j}{2}w_j^2\right\}. \tag{4.6}$$

The set of weight precisions $\alpha$, which govern the scale of the weights $w$, are drawn from a Gamma distribution, which models the distribution over non-negative precisions. In addition to non-negativity, this choice is made for several reasons. First, a Gamma hyperprior is conjugate with the Gaussian prior over the weights which will be useful when performing inference. A conjugate prior for a given likelihood function is a prior that results in a posterior distribution that is of the same family as the prior. A Gamma prior also implicitly encourages smaller weight magnitudes, leading to natural shrinkage and promoting sparsity in the inferred weights. Uninformative hyper-hyperparameters $a_0 = 10^{-2}$ and $b_0 = 10^{-4}$ which set the shape and scale of the Gamma are selected to yield a vague prior over each $\alpha_j$ defined as

$$p(\alpha_j) = \Gamma(\alpha_j; a_0, b, 0). \tag{4.7}$$

A vague or uninformative prior simply means the assumed distribution is broad. This imposes minimal assumptions regarding the parameter values, whilst still providing natural parameter shrinkage (70). This is favorable to a uniform prior which implicitly restricts parameter values and can lead to improper posteriors.

The scalar noise precision $\beta$ (inverse variance) of the residual $\epsilon$ is also modeled as a hyperparameter within the model. Given the least-squares assumption of a

Gaussian residual, we once again adopt a Gamma prior over $\beta$. The values of the shape and scale parameters are identical to those used in Equation 4.7, but are defined by the hyper-hyperparameters $c_0 = 10^{-2}$ and $d_0 = 10^{-4}$ such that

$$p(\beta) = \Gamma(\beta; c_0, d, 0). \tag{4.8}$$

**Initialization:** Models are initialized using a maximum-likelihood (ML) solution such that

$$w_{ML} = X^{\mathsf{T}} Y (X^{\mathsf{T}} X)^{-1}. \tag{4.9}$$

The ML solution is then used to initialize the residual precision hyper-parameter, $\beta$, such that:

$$\beta^{-1} = \frac{1}{N} \sum_{i=1}^{N} (y_i - w_{ML}^{\mathsf{T}} x_i)^2 \tag{4.10}$$

The ML solution provides an initial estimate that follows from the likelihood, ensuring that we start in an informative region of parameter space. This is useful for initializing the noise precision, $\beta$, and helps the variational updates from settling into poor local minima. While absolute guarantees depend on problem-specific properties, ML initialization is a well-established heuristic that improves the stability and efficiency of variational inference.

#### 4.4.2.3 Variational Inference

Fully Bayesian solutions are obtained by marginalizing over the posterior distributions of the parameters. The difficulty arises when computing the posterior distribution, which analytically is almost always intractable. Hence, sample based approaches such as Markov-Chain Monte Carlo are often employed. While these methods are good at approximating the true posterior, they scale poorly with the number of parameters included. Further, convergence is not easily assessed. Here, we adopt an approximate inference approach, called variational Bayes, referred to herein as VB. VB is a computationally tractable alternative to MCMC, bringing both scalability and a principled approach to assess whether convergence has been achieved. These attributes are particularly important for the tidal estimation procedure, as with over 300,000 individual locations, assessing convergence manually is intractable. The objective of our analysis is to infer the distributions over the individual elements of $\theta$. The basic idea of VB is to adopt analytical approximations for each distribution which can be optimized in an iterative and computationally tractable way. We first introduce an approximate posterior $q(\theta|Y)$. The functional form of this posterior is chosen to be conjugate with the prior over $\theta$ such that $q(\theta|Y)$ factorizes as

$$q(\theta|Y) = q(w|Y)q(\alpha|Y)q(\beta|Y). \tag{4.11}$$

This choice follows from the mean-field approximation, which assumes our latent parameters to be mutually independent. This selection is deliberate, as it allows for analytical marginalization over the parameter posteriors, making inference computationally efficient. However, it also neglects correlations between parameters, potentially underestimating uncertainty in the presence of strong multicollinearity. Despite this, the approach has been extensively validated for both trend extraction and tidal estimation and is thus sufficient for these tasks (105; 71).

Our objective in VB is to minimize the difference between the approximate posterior $q(\theta|Y)$, and the true posterior $p(\theta|Y)$. This difference can be assessed by considering our observable, the data evidence $p(Y)$. Using our approximate posterior, we can rewrite the log evidence $p(Y)$ as the sum of two separable terms such that

$$\log p(Y) = F\left(p(\theta|Y), q(\theta|Y)\right) + \mathrm{KL}\left(p(\theta|Y), q(\theta|Y)\right). \tag{4.12}$$

This is the fundamental equation of VB and is composed of two terms. The first term is the negative variational free energy, referred to as the evidence lower bound (ELBO). This provides a strict lower bound on the model evidence. The second term is the Kullback-Liebler (KL) divergence between the approximate and true posteriors over $\theta$. This term provides natural model shrinkage as it increases with the number of free parameters $\theta$. It can be seen that maximizing $F(p,q)$ will result in the approximate posterior being as close as possible to the true posterior. Due to the fact that $q(\theta|Y)$ can be factored as Equation 4.11, $F(p,q)$ can be maximized by iteratively optimizing each of $q(\theta|Y), q(\alpha|Y), q(\beta|Y)$ separately. Update equations for this procedure can be found in (73). An implementation of this approach can be found in the replication notebooks "Fjord_Tides.ipynb" and "seismic_attribution.ipynb".

**Bayesian R-squared** To evaluate the quality of the Bayesian regression we utilize the Bayesian $R^2$ proposed in (102). This is necessary as the variance of the predicted values can be greater than the variance of the data, thus rendering the conventional $R^2$ definition nonsensical. The modified Bayes $R^2$ is simply given by

$$\text{Bayes } R^2 = \frac{\text{Var(predicted)}}{\text{Var(predicted)} + \text{Var(residual)}}, \tag{4.13}$$

where Var(residual) is the expected variance of the errors as given by the model.

#### 4.4.2.4 Tidal Estimation

Due to the extreme sparsity of available SWOT data (less than 25 measurements over a full year), extreme care is needed when performing tidal harmonic analysis (105). Harmonic analysis assumes tides can be described by the superposition of waves at discrete tidal frequencies. These frequencies exist at harmonics of the motions between the Earth, Moon, and Sun and are described as constituents. In total, there are hundreds of possible constituents, however, for practical purposes we need only concern ourselves with a few. Given $n$ constituents, the $k^{th}$ constituent with frequency $\omega_k$ has a corresponding tidal wave of

$$C_k \cos w_k t - \phi_k = A_k \sin \omega_k t + B_k \cos \omega_k t. \tag{4.14}$$

Comparisons of tidal constituents are done in terms of the amplitude $C_k = \sqrt{A_k^2 + B_k^2}$ and phase $\phi_k = \arctan A_k/B_k$. Modern tidal analysis is carried out in the time-domain using least-squares estimation, and can thus be applied to irregularly sampled time-series. To accomplish this we define the tidal estimation problem as a general linear model, such that the observed sea-level, $y_i$, at any time is given by

$y_i = w^{\intercal} x_i + \epsilon_i$. Here, $x_i$ is the $i^{th}$ row of an $M \times N$ matrix of basis functions where $M$ is the number of measurements and $N = 2n + 1$ with $n$ equal to the number of constituents, $w$ is a set of inferred weights, and $\epsilon_i$ is the non-tidal residual. From this, we define a design matrix $X$ given by

$$X = [1, \sin \omega_0 t_i, \cos \omega_0 t_i, \ldots, \sin \omega_k t_i, \cos \omega_k t_i]^{\intercal} \tag{4.15}$$

where 1 corresponds to the bias, and the remaining values the quadrature amplitudes. The extreme sparsity of the Dickson fjord reference series is such that only the dominant lunar tide, M2, can be reliably estimated as shown in (28). Even though we are only interested in M2, it is necessary to include additional constituents in the analysis to 'soak' up the contributions of additional constituents (28; 105). After considerable experimentation, and validation against the in-situ gauge measurements, it was found that the best M2 estimation was obtained by using the semi-diurnal M2, N2, S2, and diurnal K1, and O1 constituents in the analysis. Table B2 in Appendix Bprovides the associated periods and origins of each constituent. Section 4.4.2.5, discusses validation with the in-situ CTD observations and the FES2022 tidal model.

Conventional harmonic analysis only considers measurements from a single spatial location. Due to the spatial coherence of the oceanic response to tidal forcing (17), this procedure leaves considerable information out. SWOT data provides a complete picture of the instantaneous sea-surface height throughout the Dickson fjord which can be exploited using an appropriate method. Here, we adopt the spatially coherent harmonic analysis procedure in (105). Readers are referred to (105) for a complete description of the procedure. However, the basic idea is to simultaneously estimate the quadrature amplitudes across a set of adjacent points by assuming that the amplitude at any point $P_{j,k}$ is given by the amplitude $w_{0,0}$ of the central point $P_{0,0}$ with a small offset denoted $w_{j,k}$. The linear model can be expanded to

$$Y_{j,k} = X_{0,0} w_{0,0} + \rho X_{j,k} w_{j,k} \tag{4.16}$$

where $Y_{j,k}$ is the observation at point $P_{j,k}$, $X_{0,0}$ and $X_{j,k}$ are the design matrices for points $P_{0,0}$ and $P_{j,k}$ respectively, and $\rho$ represents the probability that the observations $Y_{j,k}$ are correlated with $y_{0,0}$. The scalar $\rho$ is obtained by taking the Fischer transform of the Pearson's correlation coefficient $r$ between $Y_{0,0}$ and $Y_{j,k}$, such that $z' = 0.5[\ln(1 + r) - \ln(1 - r)]$, and $\rho = 1 - \Phi(z')$ with $\Phi = \mathcal{N}(0, 1)$. By including the probability that the observations are correlated, $\rho$, we impose an assumption that points with more similar time-series will have similar tides. Furthermore, in the limit where $Y_{j,k} = Y_{0,0}$, the quadrature amplitude $Y_{j,k} = X_{0,0}w_{0,0}$. We can further restrict the functional form of the relationship between adjacent locations to boost the ratio of data to parameters. As validated in (105), we approximate the relationship between quadrature amplitudes in adjacent locations to be linear, such that

$$Y_{j,k} = \rho \cdot (d_x \cdot X_{j,0}w_{j,0} + d_y \cdot X_{0,k}w_{0,k}) + X_{0,0}w_{0,0} \tag{4.17}$$

where $d_x$ and $d_y$ are the normalized distances between points along the horizontal and vertical direction respectively, taken to be $(-1, 0, 1)$ for convenience when using gridded SWOT data. This approach effectively doubles the ratio of data to parameters and was found to improve uncertainty estimation (105). By virtue of using gridded data with a resolution of 100 m, the assumption of linearity between neighboring quadrature amplitudes is more than sufficient.

While the spatially coherent harmonic analysis can be used in tandem with any estimator, here we make use of the variational Bayesian (VB) estimator described above in Section 4.4.2.2. This choice is based on the following reasons. First, (105) find the VB approach to be less sensitive to both stationary (Gaussian) and non-stationary noise artifacts. Second, VB provides natural parameter shrinkage, which is helpful for reducing so-called cross-talk between constituents left out of the analysis. Lastly, the implicit uncertainty information is helpful in assessing the quality of the tidal estimates. The analysis shown in Figure 4.4 only includes locations which have at least 23 observations. Figure B3 in Appendix B shows the distribution of SWOT measurements which can be used for tidal estimation through the fjord. A complete implementation of this approach, and code to replicate all tidal estimation is given in "Fjord_Tides.ipynb" (111).

#### 4.4.2.5 Tidal Validation

Tidal validation is carried out for both the SWOT derived tidal estimates, and the standard tidal estimates from FES2022 (54) using the in-situ CTD gauge as a ground truth (114). For each constituent $k$, models (VB or FES) are assessed using the root-mean-squared error, defined as

$$\mathrm{RMS}_k = \sqrt{\overline{[(A_{k,\mathrm{mod}}\sin(\omega_k t) + B_{k,\mathrm{mod}}\cos(\omega_k t)) - (A_{k,\mathrm{CTD}}\sin(\omega_k t) + B_{k,\mathrm{CTD}}\cos(\omega_k t))]^2}}, \tag{4.18}$$

where the overbar indicates the average over a complete cycle (e.g. $0 \to 2\pi$). This metric reflects the combined error in predicted sea-level height introduced by both amplitude and phase errors. We also introduce a secondary metric, the relative RMS error, which weights the RMS error relative to the magnitude of the CTD gauge constituent $\mathrm{RRMS}_k = \mathrm{RMS}_k/A_{k,\mathrm{CTD}}$. Comparisons between RMS and relative RMS error are provided in Figure B.6. For the SWOT estimates, the median amplitude and phase for the entire fjord are used. As noted above, we only include the constituents M2, N2, S2, K1, and O1 in the analysis due to the limited SWOT data. For M2 the VB model applied to the SWOT data yields significant improvements over FES2022 of 52%. The SWOT derived N2, S2, K1, and O1,in contrast, are relatively inaccurate, due to the smaller amplitudes and longer alias periods. Similar results have been found for early SWOT data by (28). In their work additional constituents were also included to yield superior M2 accuracy. When correcting the SWOT data, we utilize just the M2 from SWOT, alongside S2, N2, O1 and K1 from FES2022. As can be seen in the relative RMS panel, many of the additional FES2022 constituents exceed 20% relative RMS error within the Dickson fjord. In order to avoid introducing potential biases from incorrect tidal corrections, we do not include constituents above this 20% threshold. To combat any residual tidal variability we instead define the center of the colormap to be the midpoint of the cross-channel slope.

It is important to note that there are large seasonal changes in tides within arctic regions (115). Due to the limited SWOT data available, and the relatively short CTD gauge deployment ($\approx$ 2 months), the extent of this variability is difficult to estimate. However, since the CTD gauge was deployed during both events, validation with this gauge is appropriate as we wish to have accurate tidal estimates during this period. As more SWOT data becomes available, the accuracy of tidal estimates will improve. It is unlikely that these improved estimates will change the conclusions made by this study as we are primarily interested in the cross-channel slope, and find no evidence of this being tidally driven based on inspection of the M2 tide. However, improved tidal estimation could prove useful for understanding how the tide itself modulates the period of the seiche, and thus impacts its dynamics and dissipation.

#### 4.4.2.6 Simple Analytical Seiche

In order to estimate the total seiche amplitude, and to relate the September and October events using the observed seismic observations, we consider a theoretical seiche acting on a simplified fjord geometry. We adopt the notation and fjord geometry employed in (94) for consistency. Here, we assume the seiche to act as an oscillating horizontal force directed 160°E (perpendicular to the Dickson fjord). Due to discrepancies between the defined geometries given in (93) and (94) we avoid prescribing values for the precise fjord dimensions except where necessary. The sloshing of the seiche produces a shift of the center of mass of the body of water $x$, and can be written as

$$x(t) = \Delta x \sin \omega t. \tag{4.19}$$

Here $\Delta x$ is the amplitude of the horizontal oscillation and $\omega \approx 2\pi/92$ Hz is the frequency of oscillation. As described in (94), the amplitude of the Rayleigh waves produced is proportional to the magnitude of the horizontal force. Hence, we write $\Delta x$ in terms of the total force $F$. Taking the second derivative of the position of the center of mass we find

$$F = \Delta x \omega^2 \sin \omega t. \tag{4.20}$$

It can be seen that the maximum force occurs at the maximum displacement $\Delta x$ of the center of mass. Using SWOT data we can only observe the cross channel slope. As such, it is useful to relate the force back to the surface displacement $\Delta z$ such that

$$\Delta x = \frac{L}{3} \frac{\Delta z}{h + h_s} \tag{4.21}$$

Equation 4.20 can be rewritten in terms of the vertical displacement $\Delta z$ as

$$F = \frac{Lm\omega}{3} \frac{\Delta z}{h + h_s} \tag{4.22}$$

We recognize that the surface displacement $\Delta z = SL$ where $S$ is the cross-channel slope. Thus, the force $F$ can be written in terms of the cross-channel slope with

$$F = \frac{L^2 m\omega}{3} \frac{S}{h + h_s}. \tag{4.23}$$

Thus, we have shown that the force is directly proportional to the cross-channel slope. This allows for the direct comparison between events.

#### 4.4.2.7 In-Situ Measurements

In-situ measurements are provided by the CTD station located in the inner portion of the Dickson fjord (as shown in Figure 4.1 Panel A and can be accessed at (114)). The station provides both standard meteorological and oceanic variables. The CTD gauge is mounted on the fjord wall, inside a high density polyethylene tube to protect it from ice in winter. Here, we only make use of the wind speed and direction, density, and water depth measurements. Data are sampled at 15-minute intervals which creates severe aliasing issues for observing the 92 s VLP signal. Due to the location of the device in the inner fjord, the seiche signal magnitude decreases beneath pre-event noise levels after only a few hours and is, thus, unobservable in the data.

#### 4.4.2.8 Dickson fjord

Here we present a brief overview of the Dickson fjord. A more comprehensive description of the physiography and climate of the fjord system can be found in (94). The Dickson fjord sits at the terminus of the Hissinger Glacier in the northernmost area of the Kong Oscar fjord system situated in East Greenland (See Appendix B Figure B1). The fjord itself sits deep in the Arctic Circle, and is thus characterized by sea-ice over much of the year. Sea ice dissipates in July and then forms again in October. The fjord fills a U-shaped valley basin, with multiple smaller glaciers situated on each side. The fjord itself is 38 km long and between 2.5-3.2 km wide. The depth ranges from 150-200 m to 700 m from West to East with an approximate depth of 540 m in the center of the fjord across from the landslide location. Bathymetry estimates are taken from a 2018 survey by the National Danish Hydrographic Office at a resolution of 15 m. We note that no data exists between 150-300 m meters of the coast due to the limitations of the survey vessel. This missing data creates large uncertainties in these regions which can significantly influence numerical simulations.

#### 4.4.2.9 Tsunami Information

Both tsunamis originated from landslides occurring in the same gully situated beneath an unnamed glacier (94). These landslides were caused by debuttressing of the glacier following glacial thinning over the past decade. Direct observation of landslide scarring, and dirtying of the glacier using satellite imagery by both (94; 93) confirm this theory. Additionally, (94) evaluate the landslide dynamics of the September 16th landslide via seismic inversion. While the October 11th event did not produce a new landslide scar, a Sentinel-2 image showed significantly more erosion than was present after the September 16th event.

Empirical evidence of the two tsunamis is given by a combination of nearby in-situ sensors, and observed run-up height. Using satellite imagery, both (94; 93) observe an initially 200 m run-up height near the location of the slide, with an average of 60 m run-ups being observed through the remainder of the fjord for the September 16th event. Tsunami run-up for the October 11th event was only observed 200 m west of the gully and was approximately 75% of the magnitude of the September event in this location (60 m vs 80 m). Almost 72 km away at the Ella Ø research station, the initial run-up height was in excess of 4 m creating significant local damage. The location of Ella Ø relative to the Dickson fjord is shown in Supplementary Figure 1. To our knowledge, no information exists regarding the run-up at Ella Ø for the October 11th event due to arctic winter darkness.

#### 4.4.2.10 Dissipation the Seiche

While the following discussion does not impact the conclusions made in this thesis regarding the initial size of the seiche, I wish to highlight some interesting aspects of the seiche dissipation. One of the most peculiar aspects of this event was the persistence of the seiche for over a week for both events. There are several unique characteristics which gave rise to this behavior. First, the energetics of the initial tsunamis which produced these seiches were massive (94). As computed, this amounts to initial seiches of approximately 7.9 m and 4 m respectively. The Fjord geometry itself also contributed to the duration of these events. Inspection of the SWOT observations shows the 'jack-knife' turn at the entrance to the fjord contains the seiche and thus slows its dissipation into the open ocean. Additionally, the near rectangular shape of the fjord body seems to have produced a singular dominant mode of oscillation, rather than many smaller modes. When combined with the fact that the fjord is relatively deep (approx 500 m), the impact of bottom friction is much less than in shallow lakes leading to the long duration. As suggested by other authors (94; 93), the eventual dissipation of the seiche was likely driven by the combination of bottom drag, internal wave radiation into the open-ocean, and turbulent mixing.

Another interesting aspect of the seiche dissipation was the observed modulation of the seiche period. The authors of (93) observed an approximately 6 hour modulation of the seiche frequency ranging from 10.874 to 10.879 mHz which they attribute to tidal influences (their supplementary Figure S1 Panel C). They suggest a possible mechanism is the change of depth modulating the speed of the wave $c$ through the relationship $c = \sqrt{gH}$ where $g$ is the gravitational acceleration and $H$ is the depth. However, a simple inspection of both the September and October VLP signals shows that the seiche actually behaves in the opposite manner (e.g. longer periods at higher tides).

We here consider an alternative mechanism: stratification. Traditionally, to leading order, the surface seiche can be treated as decoupled from the internal seiche because the former behaves as a gravity wave traveling at speed $\sqrt{gH}$, whereas the internal seiche propagates more slowly at $\sqrt{gH\Delta\rho/\rho}$, where $g$ is gravitational acceleration, $H$ is depth, $\rho$ is the mean water density, and $\Delta\rho = \rho_2 - \rho_1$ is the density contrast from bottom to top. The decoupling is justified by the typically small value of $\Delta\rho/\rho$, leading to vastly different wave speeds. However, in stratified systems with non-horizontal bathymetry or open boundaries, surface seiches can generate internal waves, which in turn can feed back onto the surface motion and modify its properties (116). In particular, (116) demonstrated that in semi-enclosed domains connected to the deep ocean—such as the fjord studied here—the surface seiche can also exhibit stratification dependent propagation given by $\sqrt{gH\Delta\rho/\rho}$. This baroclinic-wave drag effectively acts as an energy sink, causing damping of surface seiche—often exceeding bottom friction. Moreover, this outward radiation modifies the boundary condition at the fjord's mouth, shifting the resonant frequency of the fundamental seiche mode. Inspection of the correlation between surface-water elevation and density $\rho$ shows a strong positive correlation (Appendix B Figure B7). Based on this observation we proposed the following explanation. At low-tide, the water column is strongly stratified with fresh (less dense) glacial runoff sitting on top, and denser salty water sitting beneath $\Delta\rho > 1$. As the tide comes in, the water column mixes, leading to the observed higher density $\rho$ at the surface and thus less stratification $\Delta\rho < 1$. Due to the fact that the fjord is quite deep $\approx 500$ m, the impact of $\Delta\rho \gg \Delta H$ as observed. Due to the fact that the CTD gauge only provides observations at the surface, it is impossible to conclude definitively that the density fluctuations observed are indicative of the stratification between $\rho_1$ and $\rho_2$. Hence, further simulated study is needed to validate this hypothesis.

### 4.4.3 Data Availability

All data used in the study have been deposited in a Zenodo repository alongside code to process and produce each figure `https://doi.org/10.5281/zenodo.15166491` (111). The L2 HR SWOT data can be accessed freely through PO.DAAC `https://doi.org/10.5067/SWOT-PIXC-2.0` (117). All code needed to replicate the given analysis can be found in the zenodo repository `https://doi.org/10.5281/zenodo.15166491` as well as the dedicated code ocean repository `(https://codeocean.com/capsule/5272497/tree)`.

# 5 Automating the tidal response method

# Abstract

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

## 5.1 Introduction

The analysis and prediction of water levels driven by astronomical and other phenomena is important for countless scientific and engineering applications. As tides can become distorted in coastal and estuarine environments, we here define these processes by the forcing which gives rise to them. Tides are the weakly nonlinear response of sea-levels to the combined gravitational and radiational[1] forcing from the moon and Sun. Tidally driven phenomena by extension are oceanic processes which exhibit coupled and weakly nonlinear responses to both tidal forcing and other external forces. Prediction of these processes is done using both numerical and empirical approaches which will be discussed below. In contrast, tools which can also be employed for the analysis of these phenomena are largely limited to tidal harmonic analysis (HA) (10) with notable exceptions for particular phenomena such as tidal rivers (118). While in most cases HA is sufficient for the study of tides, tidally driven phenomena such as tidal rivers and storm surges often violate the assumptions implicit within this formulation (105). Application of HA in these contexts leads to biased estimates of tidal constituents and limited predictive power of the coupled process (119). As will be shown, it is inappropriate to apply HA in these contexts, yet no general purpose method exists for both the analysis and prediction of these processes. Empirical approaches rely on concurrent observational data and thus struggle if such data is unavailable (e.g. predictions beyond short-term or if gauges are decommissioned). This manuscript addresses this deficiency by automating the tidal response method devised by Munk and Cartwright in the 1960s (5). The resultant method enables empirical analysis and prediction of any tidally driven process without requiring concurrent water-level measurements. To highlight the flexibility of our approach, all case studies in this manuscript were run using just three lines of code, without modification.

---

[1] Radiational forcing refers to non-gravitational forcing such as surface pressure fluctuations from atmospheric tides.

### 5.1.1 Literature Review

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

The standard approach to predicting past and future tides, based on sea-level observations, is through harmonic analysis (HA) (10). As proposed by Darwin and Kelvin (120), with refinements from Doodson (11), HA describes the tides as the sum of oscillating tidal constituents whose amplitudes and phases are typically determined through least-squares fitting to the observed sea-level. The linear constituents described by HA exist at nearly identical frequencies to the astronomical forcing (47). Shallow water effects in HA are handled by including nonlinear constituents whose frequencies are the sums or differences of linear constituents. This approach imposes strict assumptions about the data. In particular, conventional harmonic analysis assumes the signal to be stationary and that the oceanic response is coherent with the astronomical forcing which gives rise to it (10). This characteristic becomes problematic when HA is applied to observations contaminated with non-tidal forcing which can result in significant biasing of the derived constituents (119). In addition to being a large source of error for operational surge forecasting, biased tidal estimates limit our ability to study both non-tidal and tidally driven processes as will be shown in the three subsequent case studies. Furthermore, evaluating tides within a harmonic framework limits the physical conclusions that can be drawn about the dynamics of the interactions of tides with non-gravitational forcing.

Machine learning methods have shown great promise in improving forecast accuracy over both conventional numerical and empirical approaches across many scientific applications (121; 122; 123). Due to the great success of HA, applications of ML to pure tidal height forecasting are limited. The nonlinearity of tidal currents and their importance for tidal energy has justified the development of several ML approaches for tidal current forecasting confined to mainly short-term predictions (20; 124; 125). Several works have developed ML approaches for storm surge forecasting (126; 127; 128). Many approaches have been tried; however, the vast majority conform to a similar approach: based on water-level observations, have a model learn how to predict the next time-step recursively based on forecasted meteorological conditions. The variation in these methods comes from both different choices of input data and preprocessing as well as various model architectures (128). While these works demonstrate compelling reductions in error compared to standard numerical models, none are used operationally (128). This is a consequence of two main shortcomings. First, reliance on in-situ observations is problematic as gauges can go offline, or have variable accuracy due to bio-fouling and environmental conditions. The accuracy of the final model is thus fundamentally intertwined with the operational conditions of the gauge used to train the model. Furthermore, these approaches can not be used where observational data no longer exist. This is problematic as it prevents the application of these methods in ungauged regions (see Section 5.4.1). Second, as a consequence of their autoregressive and black-box construction, these methods have limited physical interpretability (128). This stems from working from observations rather than forcing. Here, the model is learning statistical correlations between these observables which describe their evolution. It is thus difficult to attribute the role of each input in producing a given forecast and model generalization is limited by past observed conditions.

Several specialized approaches have been developed for the analysis and prediction of nonstationary sea-levels. First, the familiar harmonic analysis formalism has been adapted for the prediction of riverine tides by replacing the conventionally static amplitudes and phases with functions of the river outflow and the diurnal tidal range from an adjacent location free from fluvial influence (118). This approach, termed "NS_Tide", derives its basis functions from a simplified fluvial model, and uses the regression to iteratively optimize coefficients for deviations in theory. This approach faces two key shortcomings; it requires (i) data from an adjacent reference station free from fluvial influence, and (ii) it has only been developed for fluvial (using river outflow) and atmospheric pressure forcing (118). As such, even for tidal rivers, the approach loses applicability if either an adjacent reference station or river outflow data is not available. Section 5.4.1 considers an application where both of these things are true. Furthermore, while “NS_Tide” can be effective for routine analysis and prediction of tidal rivers, it cannot be viewed as a general tool for analysis and prediction of nonstationary tidal processes. Other approaches, such as the species concordance method, similarly rely on uncontaminated reference series and thus face similar limitations (35). Many approaches have been developed for analysis of tidal processes, but offer no predictive value and are thus out of the scope of this discussion (129; 130; 45).

The complex modifications of HA to accommodate nonstationary effects are an artifice arising from restricting the solutions to sums of discrete harmonic functions of time. In contrast, the Response method provides a framework for computing the weakly nonlinear oceanic response to a set of input forcing (5). This approach can be viewed as learning a time-invariant operator from the forcing to the associated sea-level (described at length in Section 5.2.1). Originally, this was done just from the tidal input potential but was later applied to meteorological forcing (16). Under this construction, the oceanic response is explicitly separated out from the input forcing, providing the ability to attribute the contribution from each input force. The learned model can therefore be used to identify the interaction of tides with these forces and can be used to predict the combined response of any tidally driven phenomenon. Many works have found response approaches to yield modest improvements over HA for pure tidal prediction (5; 39). However, applications of the approach to predict nonstationary tidally driven phenomena have been limited outside of (16). This is a consequence of the approach being challenging to automate as nonstationary tidal processes exhibit irregular and nonlinear response functions (see Section 5.2.1.1) (5; 16; 131). The remainder of this chapter addresses this limitation and develops a general-purpose analysis and prediction tool for tidal processes.

## 5.2 Methods

### 5.2.1 Response Method

The Response method seeks to describe the time-invariant oceanic response at a point to a series of impulses $V(t)$. The time-series representing the tidal input potential $V(t)$, can be computed at any time and location based on the known orbits of the Moon, and Sun. A key contribution of the original response method was the recognition that the tide at a given point was not merely driven by the instantaneous equilibrium potential, but rather a combination of its past, present, and future values (5). Hence, the Response method expands the gravitational tidal potential using spherical harmonics such that

$$V(\theta, \lambda; t) = g \sum_{n=0}^{n} \sum_{m=0}^{m} [a_n^m(t) U_n^m(\theta, \lambda) + b_n^m(t) V_n^m(\theta, \lambda)], \tag{5.1}$$

where $g$ is the gravitational constant, $\theta$ is the Geographical Colatitude, $\lambda$ the East longitude, $U_n^m$ and $iV_n^m$ are complex spherical harmonics of degree $n$ and order $m$, and $a_n^m$ and $b_n^m$ are the associated complex values of the global tide function $c_n^m(t)$, yielding $V(\theta, \lambda; t) \in C$. The spherical harmonics converge rapidly and are defined so that the tidal species of interest is governed by the selection of $m$. Normalization of the spherical harmonics is consistent with quantum mechanics (132). A table containing the frequency bands for each species is provided in (5).

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

$$\hat{\zeta}(t) = \sum_{m,n} \sum_{s} [u_n^m(s) a_n^m(t - \tau_s) + v_n^m(s) b_n^m(t - \tau_s)], \tag{5.2}$$

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

$$x^{th}\text{order response} = \sum_i \cdots \sum_x \sum_s \cdots \sum_{s'} w(i, \ldots, x, s, \ldots, s')(c(t-\tau_s))(\ldots)(\ldots)(c(t-\tau_{s'})), \tag{5.3}$$

with the number of terms increasing exponentially with the order of the interactions. Typically, higher order terms are limited to only one lag for computational efficiency. Due to the lack of orthogonality of the basis functions for the standard response method, much care is required when performing a sequential analysis. The time lag, $\tau_s$, is typically 2 days and $s$ is the maximum lag with the understanding that convolutions are carried out using $\tau = 0, \Delta\tau, 2\Delta\tau, ..., s\Delta\tau$. An extended discussion of the choice of time lags and their implications on the associated response is given in Section 5.2.2.3.

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

A complete description of the Gravitational and Radiational input functions can be found in their original paper (5). Additionally, code to compute both gravitational and radiational input functions is provided in the RTide[2] python package.

#### 5.2.1.1 Challenges of Conventional Response Analysis

While the response method possesses several advantages over HA, it carries with it what is referred to as the burden of realistic input function selection (5). This language is in someways a misnomer as the input forcing functions for a standard response analysis are readily computed from the derivations given in (5). We argue the difficulty in a response analysis is not introduced by the input functions themselves, but instead, the need to specify the analytical form of their interactions *a priori*. Here, we refer to these relationships as *interaction kernels*, as this description is consistent with the broader literature on Volterra series. Similar to the selection of harmonic constituents in the harmonic method, the standard response approach requires the definition of these interaction kernels beforehand. This requirement is in some ways more restrictive than the selection of constituents as these interactions becomes increasingly complex and consequently difficult to prescribe in strongly nonlinear regions. Further difficulty is introduced if one wishes to include other forcing inputs as is shown in (16). Indeed, as will be shown in 5.4.1, the nature of these interactions for many multivariate inputs is not known. This is made more difficult to debug as the response weights generally lose meaning for such inputs (16). As such, the standard response approach often requires expert trial and error which has severely limited its automated application.

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

---

[2] `https://github.com/thomasmonahan/RTide`

### 5.2.2 ML Response Framework

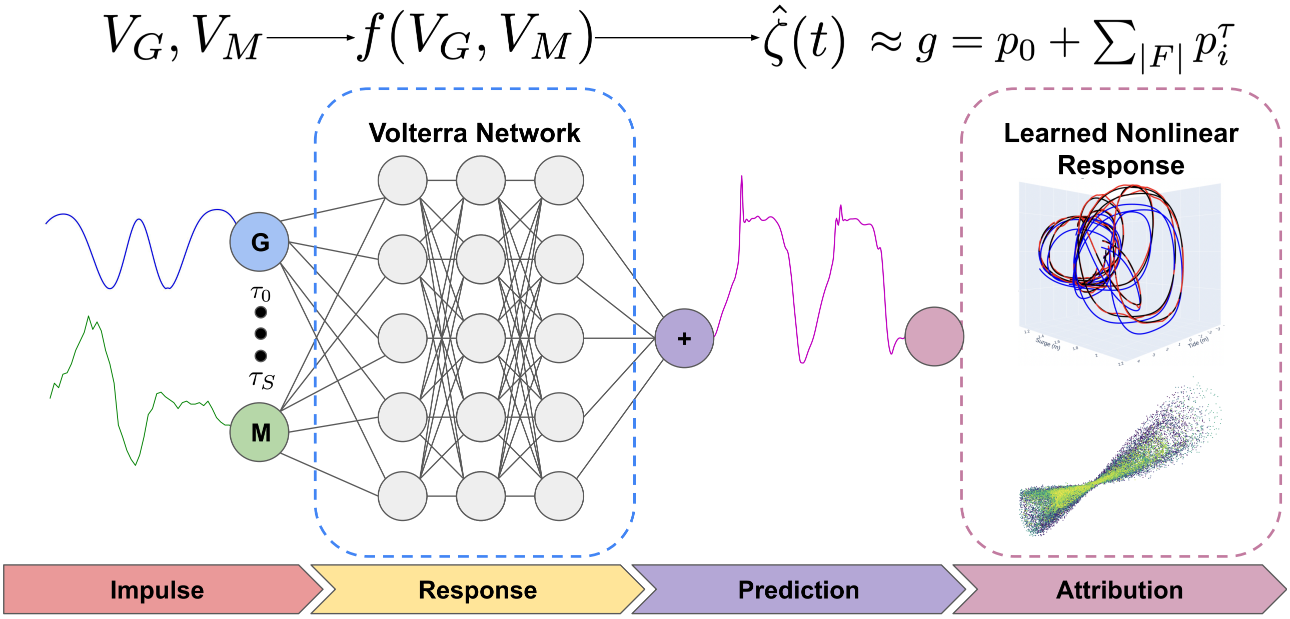

Figure 5.1: **Schematic of the proposed ML Response framework.** The input consists of the lagged Gravitational ($V_G$) and Multivariate ($V_M$) input functions. Here multivariate can be anything from the 'radiative flux' proposed in (5) to river levels or atmospheric variables. The oceanic response function $f(V_G, V_M)$, equivalent to Equations 5.2 and 5.3, is learned by a three-layer neural network (Volterra network) and used to predict either total sea-level $\hat{\zeta}$ or the astronomical tide $\hat{\zeta}_T$ at any time instance $t$. Learned models can then be taken apart and studied by transforming into an equivalent linear model $g$ using SHapley Additive feature exPlainers which quantify the contribution of each feature $p_i^{\tau}$ to the final prediction (6) outlined in Section 5.2.3.

In this section, we introduce our proposed ML Response Framework. The basis of our approach is the integration of a specific class of three-layer neural networks into the tidal impulse-response theory. That is, rather than specify a parametric form for the interaction kernels before an analysis and infer the response weights using least-squares, as in the original method, we instead use the networks to infer both the interactions between inputs and the associated weights simultaneously and directly from the data. The motivation for this procedure lies in the provable equivalence between the Volterra series—the mathematical basis of Munk and Cartwright's response method—and the aforementioned constrained three-layer-perceptrons (133). This relationship makes the approximation of any arbitrary combination of Volterra kernels possible (133). Thus, our framework can infer the appropriate Volterra kernel structure directly from the available data and efficiently represent higher order nonlinearities (e.g. 3rd order and above) if required (134).

The response prediction, defined by Equation 5.2, can be replaced by

$$\hat{\zeta}(t) = f(V_0(t-\tau_0), V_0(t-\tau_1), \ldots, V_0(t-\tau_S), \ldots, V_X(t-\tau_0), V_X(t-\tau_1), \ldots, V_X(t-\tau_S)), \tag{5.4}$$

where $f(\cdot)$ is the NN approximation of the time-invariant impulse-response function and $V_i$ are the set of all input functions at the associated time lags. It should be noted that for input functions expanded in spherical harmonics, inputs of differing degrees and orders are provided as separate inputs, and thus considered separate input functions, as in the original formalism. In contrast with the original method, the data does not need to be bandpassed beforehand as all species and interactions are inferred simultaneously. The underlying mathematical connection to the Volterra series allows the networks, as deployed herein, to simultaneously learn both the linear time invariant impulse-response, as well as the associated nonlinear interactions, which are challenging for the original rigid convolution method. For the analysis presented in this paper, the network weights $w$ are computed using gradient descent. This non-parametric procedure retains the physical advantages and realism of the original response method but overcomes the limitations of the parametric original method by allowing the model to simultaneously learn the structure of the input function interactions and weights directly from the data. The simultaneous inference of feature weights and interactions serves as a natural regularizer, thereby avoiding some of the variance-bias trade-off imposed by least squares. Critically, when paired with the training and hyperparameter optimization procedures presented in Sections 5.2.2.5 and 5.2.2.4, this framing reduces the need for expert tuning and relaxes the need for the explicit definition and expansion of input functions. The following section outlines the implementation of the method used in this work and in the RTide Python package, and describes its relation to the original response method. Figure 5.1 provides a schematic of our approach.

#### 5.2.2.1 Network Architecture

We use a class of three layer feed forward neural networks (NN) termed "Volterra Neural Networks" due to their ability to represent any arbitrary Volterra series (135). The basic NN architecture has three layers with size equal to the number of inputs. If we have $X$ input functions and $S$ time-lags then the size of each layer is simply $X \times S$. We find little sensitivity to the selection of activation functions, but utilize the sigmoidal activation function, tanh, which allows for the indirect estimation of Volterra kernels from the learned network (135). The final layer consists of a simple adder which sums the outputs directly from the previous layer. The choice of this architecture reflects our objective to produce a non-parametric approach that is mathematically equivalent to the original method, but critically removes the need to select interaction terms *a priori*. As a consequence of this equivalence, these networks are generally robust and can be applied universally. In our testing (results not included in this manuscript), we have found that structuring networks in accordance with known physical interactions between input functions (e.g. tide-surge interactions) can improve the resolution of the underlying dynamics. We believe the development of such inductive biases to be a promising avenue for future work.

A key modification in our framework is the use of a "deep" rather than "wide" network. A wide network is comprised of a single hidden layer and is equivalent to a generalized linear model. This construction yields excellent memorization capabilities, but will often struggle with generalization (136). As such, it is easy to overfit these models which will lead to poor performance on unseen data. Conversely, deep networks, with multiple hidden layers, demonstrate superior generalization as they can form low-dimensional embeddings from the input space. A wide model is the appropriate choice when applied to low-noise data, and in predominantly linear tidal regimes. However, we find that the improved generalization of deep networks is beneficial for studying complex tidal regions with non-tidal forcing.

Typical training times are approximately 3 minutes when applied to 1-year of hourly data (8760 training points) using $X = 10$ input functions with $S = 7$ time-lags ($3\times(10\times7)^2+3\times10\times7+10\times7+1 = 14981$ trainable parameters) using a Tesla K80 GPU. While this reflects the total number of trainable parameters, the regularization of the model reduces the effective number of degrees of freedom drastically. As such, the data helps to inform the complexity of the model which we believe to be a major advantage. We note that training time scales with both the number of measurements and the number of input functions and time lags included.

#### 5.2.2.2 Input Forcing

This work defines standard analysis as any analysis which makes use of just the gravitational input functions defined in (5). These are implemented exactly as in the original work up to degree 3 for the gravitational inputs. Due to the self-regularization of our non-parametric approach, it was found useful to include all available degrees and orders of spherical harmonics for standard analysis. Thus, when combined with hyper parameter optimization, standard analysis can be fully automated.

As noted in (5), inclusion of the 'radiational flux' can help combat the variable nature of seasonal and annual weather fluctuations. Depending on the location, and the consistency of this seasonal variability, inclusion of such inputs can either help or hurt long-term prediction skill. For Sections 5.3.1 and 5.4.1 we do not utilize radiational input functions as we wish to isolate and obtain physically meaningful insights regarding the oceanic response to different forcing. Section 5.4.2 briefly evaluates the impact of including Radiational inputs in Table C.1. Radiational inputs can be included in the RTide package by setting 'radiational= True' in the 'Prepare_Inputs' function. Due to their quasi-physical nature, the learned response should also be regarded as quasi-physical.

The flexibility of our framework enables the inclusion of any multivariate input forcing which can be represented as a time-series in addition to the original gravitational and radiational input functions given in (5). In this work, we also include what we term multivariate input functions. We define multivariate inputs as any forcing that is not gravitational. A consequence of our non-parametric framework is that we can relax the requirement for input function definition. As an example, in (16), the pressure gradient is obtained by fitting coefficients of a spatial Taylor series to meteorological observations at 8 reference stations. Our testing showed that our method could obtain similar performance when trained using the meteorological variables as direct input. In this way, our non-parametric formulation does not impose approximate physical assumptions such as the linearization of wind stress used in (16). Rather, our approach infers these relationships directly from the data. This characteristic can thus be used to inform the design of such low order models and is described in Section 5.4.1. This capability makes the testing of new input functions extremely easy and the integration of other data streams trivial using only one line of code in RTide.

#### 5.2.2.3 Lag Considerations and the Credo of Smoothness

The selection of which time lags and spacing to use, often referred to as *embedding parameters*, is critical to the success of a Response analysis. In their original paper, Munk and Cartwright propose a uniform (evenly spaced) time lag of $\Delta\tau = 2$ days based on empirical analysis of the Honolulu tidal record (5). Here, we refer to the selection of evenly spaced lags as uniform embedding. They further truncate their response analysis after $S = 3$ lags following the observations that the tidal admittance $Z(f)$ shows no indication of "wiggles" with period less than 1/6 cycles per day. This selection of a maximum lag results in a smoothing of any portion of the signal with period less than 1/6 cycles per day which they termed the "credo of smoothness". While our testing confirms these observations, we also found that the inclusion of additional non-uniformly spaced lags improved performance (See Table C.1 in Appendix C). We remark that while this approach has improved performance for our non-parametric method, the non-uniform lags may introduce complications when performing spectral analysis on the classic method.

We believe the justification for this modification comes from the dynamics of the tidal system. It is well understood that uniform embedding can yield poor results when applied to systems with multiple strong periodicities across disparate timescales (137). Non-uniform embedding provides the opportunity to characterize systems with fast-slow dynamics (138). While numerous data-driven approaches have been developed for the automated selection of non-uniform time embeddings, tidal flows provide a unique circumstance in which the characteristic periods are well understood. We utilize a non-uniform embedding scheme based on a subset of the natural tidal frequencies. The associated time lags $\tau_i$ are given by the quarter period heuristic which yields a maximally convex embedding in the reconstruction space if the characteristic frequencies can be appropriately identified (139). We select $S$ positive and negative time lags $\tau_i$ such that

$$\tau_i = \frac{2\pi}{4\omega_i} \tag{5.5}$$

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

As proposed in (49), we have experimented with a symmetrical analysis scheme that includes an equal number of positive and negative lags. We found this procedure to improve model performance (See Table C1 in Appendix C). While leading time lags appear to violate causality, we propose the following explanation based on the method of characteristics. As a consequence of selecting a finite number of lags in our embedding procedure, we constrain our lag embedding to consider the first $\tau_i$ negative lags. Suppose we have an arbitrary signal whose characteristic period is not included in this procedure. If we consider the tidal signal to be a series of waves observed at a stationary location at time $t$ governed by a single PDE, then the signals ‘originating’ from each time lag can be considered to act along the Spatiotemporal-Characteristic of the tidal system. The inclusion of positive lags, therefore has a simple explanation; positive lags correspond to signals acting along the same characteristic, but intersecting a $\tau_i$ in the future (whether in time or space).

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

#### 5.2.2.4 Hyperparameters and Tuning

While the selection of input functions and interaction kernels has now been automated, the choice of hyperparameters plays an important role in the success of an RTide analysis. Here, hyperparameters refer to the parameters which govern the training characteristics of the model (training duration or ‘epochs’, learning rate, and regularization strength). Reported values as well as general recommendations for hyperparameters are included in the RTide package. In our testing, we found the default hyperparameters included in the RTide package to yield state-of-the-art performance and generalization (balanced performance out of sample) across a broad range of problems of varying complexity and non-gravitational forcing. Indeed all case studies presented here have been produced using these standard parameters. However, to allow users to achieve maximal performance an adaptive hyperparameter selection scheme can be adopted. We elect to use the Keras Hyperband tuner due to its rapid convergence and strong performance NN architectures (140). The optimization procedure is carried out in a tournament style and is accelerated through optimal resource allocation and early stopping (141). This procedure is carried out during the training procedure and repeated until a user defined threshold (either iterations, time, convergence, etc.) is met.

#### 5.2.2.5 Training Procedure

For problems where data is abundant (e.g. ≥ 30 days of data with sampling rate ≤ 1 day), a standard training procedure is sufficient. That is, after weight initialization, the entire network is trained end-to-end using the computed input functions at the associated time lags as direct input. The input training time-series is randomly shuffled and automatically partitioned into train and validation sets. By default, this is set to 95% and 5% of the data respectively. This validation set is used throughout the network training to monitor the model's generalization capabilities. Inspection of the training set and validation set loss curves is helpful for diagnosing overfitting and is shown by default when performing RTide analysis. An example of this is shown in Figure C.2 in Appendix C.

Unlike conventional response analysis in which the interaction terms must be defined beforehand, the interactions between inputs and associated weights are simultaneously inferred through the training procedure using gradient descent. This allows the data to determine the structure of the model thereby acting as implicit regularization. Using the RTide implementation this allows for response analysis to be carried out using only three lines of code. Gradient descent simply looks to minimize the provided cost function. Hence, additional modifications could be applied to the loss function depending on the user's objectives. For example, minimizing the absolute error in water-level prediction will not always yield the best tidal estimate and vice versa. The results reported in this work only use the mean absolute error (MAE) as the chosen loss function. We note that early testing has shown the Surface Similarity Parameter (142) can yield more accurate tidal estimates in some scenarios and is implemented in RTide. Unlike Mean Squared Error, MAE places less emphasis on outliers and is thus more robust to nonstationary artifacts. Dynamic tuning of the learning rate by the ADAM optimizer further reduces sensitivity (143). We have also experimented with both L1 and L2 regularization of the network weights. Both of these have proved to be beneficial for reducing overfitting and reducing the effective number of degrees of freedom of the model. Regularization strength parameters are included in the hyperparameter tuning described in the previous section.

### 5.2.3 Model Interpretability – SHAP Feature Importance

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

Shapley Additive exPlanations (SHAP) is a model agnostic framework for developing machine learning 'explainer models' based on Shapley values from game theory (6). This framework allows us to transform the learned response model $f$ into a simplified additive explanation model $g$ defined by the linear combinations of binary variables $p_k$ such that $f \approx g = p_0 + \sum_{|F|} p_k$ where $F$ is the set of all input features and $|F|$ is the number of input features[3]. For a given prediction $f(t)$, the associated Shapley values $p_i^\tau(t)$ for each input function $i$ and time lag $\tau$ provide an allocation of credit among input features (e.g. the sum of SHAP values is equal to the expected value of the model) (144). The Shapley value $p_i^\tau(t)$ of each feature can therefore be interpreted as the impact of each input function at the associated time lag on the oceanic response.

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

$$p_i^\tau(t) = \sum_{R \in F/\{^\tau_i\}} \frac{|R|!(|F| - |R| - 1)!}{|F|!} [f_{R\cup\{^\tau_i\}}(t_{R\cup\{^\tau_i\}}) - f(t_R)] \tag{5.6}$$

---

[3] A *feature* is the machine learning terminology for a single input to the model. In our case features correspond to an individual input function, at the associated time-lag $\tau$.

Shapley values satisfy two important properties:

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

The combination of these properties makes it possible to explain why the learned response model makes its prediction and thus attributes physical causes of variability. Critically this analysis also allows for the comparison of feature importance between locations, times, and in the presence of features with partial coherence. Indeed, implicit within this is quantifying the interactions of different inputs as shown in Section 5.4.1. This is important as it can shed light on how non-gravitational forcing modulates the astronomical tide. A complete treatment of the SHAP algorithm including proofs of the two aforementioned properties can be found in (6). SHAP values in this work are computed using the SHAP python package (github.com/shap/shap).

### 5.2.4 Harmonic Analysis

The Harmonic Analysis described in this paper uses a Python adaptation of the MATLAB Harmonic Analysis program UTide (34) with a Rayleigh threshold $R = 1$. Analysis is conducted using ordinary least squares (OLS) unless explicitly stated. A complete description of the harmonic analysis theory can also be found in (10).

## 5.3 Simulated Results

### 5.3.1 Nonlinear Impulse Response Discovery

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

$$\zeta = \zeta_T + \zeta_S + B * \zeta_S(\zeta_T'' + \zeta_T'), \tag{5.7}$$

where $\zeta_T$ and $\zeta_S$ are the independent tide and surge signals, with tide-surge interaction defined by a constant $B$, and the first and second derivatives of the tidal signal $\zeta_T'$ and $\zeta_T''$. The simulated tide $\zeta_T$, sampled every 15 minutes (70080 observations), is composed of 59 harmonics originally derived from 1-year of tide-gauge measurements at Eastport, Maine. The signal is largely semi-diurnal with an amplitude of approximately 3 meters. We approximate a surge using the formula outlined in Cartwright's original work on surges (16) such that $\zeta_S = A * (1 + \cos(2\pi t/l))$ where $6 \leq l \leq 96$ (hours) is the duration of the surge, defined from $0 \leq t \leq l$ and $0.1 \leq A \leq 2.1$ (m) is the amplitude of the surge. The nonlinear tide-surge interaction term is defined by the product of the pure surge with the first and second derivatives of the tidal signal respectively. This simplified model can approximate the well documented behavior of tide-surge interactions in which surge effects are amplified when they coincide with a low and rising tide and damped at high and falling tides (145; 146; 147). It should be noted that this is a low-order approximation developed from observations in order to validate synthetic models of tide-surge interaction and is not a true representation of the physics. $B = 1$ is chosen such that the tide-surge interaction never exceeds more than 50% of the actual surge (148). We elect to use this simplified tide-surge formulation as it is easily interpretable within the time domain, and illustrates the importance of our approach even if the surge signal was perfectly predicted. Simulated surges are randomly generated using the defined parameter bounds and randomly added to the synthetic tidal signal so that a surge event occurs within each 2-week window. In total, there are 52 surges of variable magnitude and duration applied to the tidal signal. White Gaussian Noise is added with a SNR of 60 to represent more realistic conditions based on analysis of the original Eastport tide-gauge data. Visualizations of the simulated surge data can be found in Apprndix C Figure C1.

Here, we suppose that the surge is produced from fluctuations in atmospheric pressure through the inverse barometer effect such that $\zeta_S = -0.01 \times P$ where $\zeta_S$ is in meters and $P$ is in mbar (149). In this case-study, let us also assume the atmospheric pressure is accurately forecasted and can thus be used as an input function into RTide. Hence, both the surge and the nonlinear relationship between tide and surge is unknown *a priori*. Figure 5.2 compares the performance of RTide and HA when used to recover the tidal signal, storm surge (inverse barometric pressure effect), and nonlinear tide-surge interaction. Both models are trained on the first 1.5 years (550 days) of synthetic data.

Before considering any non-tidal effects, we first evaluate the performance of conventional least-squares (OLS) HA with RTide. Figure 5.2 Panel (A) shows that the process of simultaneous tide, surge, and tide-surge estimation results in the more accurate recovery of the true tidal signal by RTide. The RTide predicted tidal signal is simply obtained by passing a vector of zeros for the normalized barometric pressure input functions. This improvement is reflected by the reduction in mean absolute error (MAE) in the derived tidal signal from 0.054 m for HA to 0.009 m for RTide. We note that this comparison is conducted using HA with OLS, however, we have also compared the performance to modified HA using iteratively reweighted

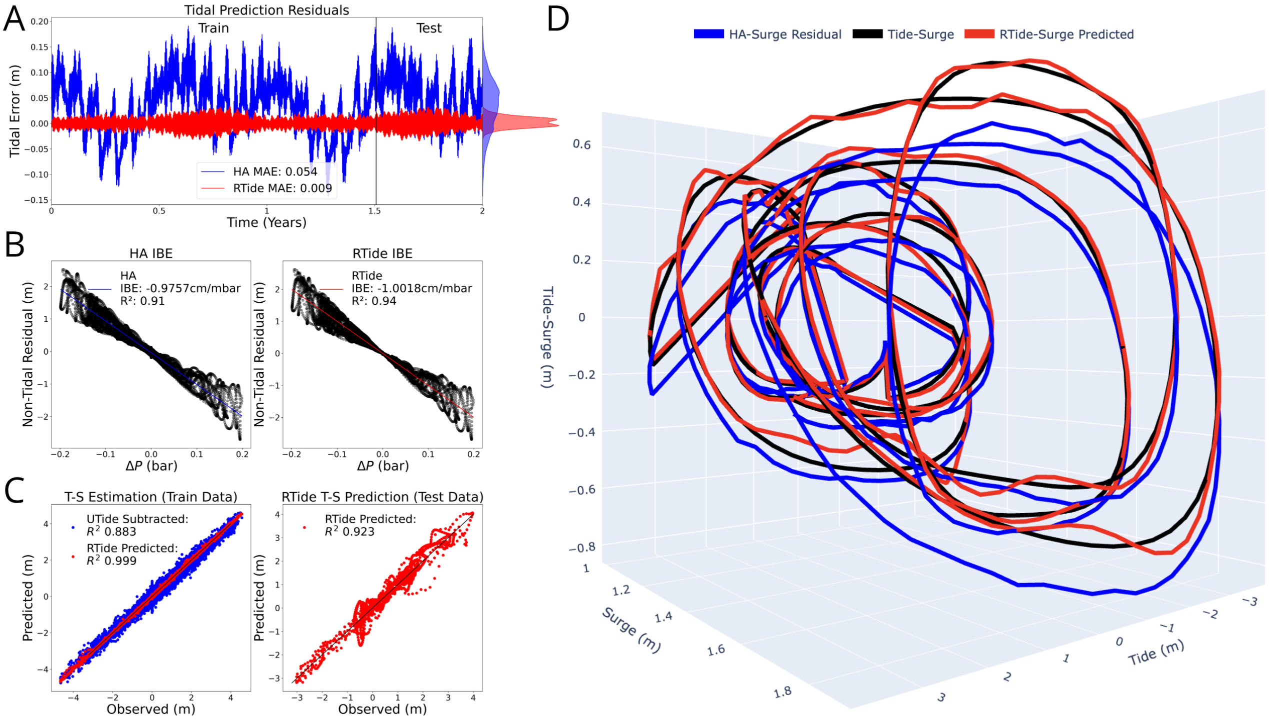

Figure 5.2: **Comparisons of RTide and HA when used to recover both the nonlinear tide-surge interaction and the tidal signal.** Panel **(A)** shows the error in tidal predictions for HA and RTide when compared to the true tidal signal. Panel **(B)** shows the HA and RTide estimated inverse barometer effect (IBE) using a linear regression of the change in pressure $\Delta P$ and non-tidal residual. Panel **(C)** shows comparisons of the HA tide-surge residual and the RTide tide-surge prediction. Note, since the HA is only a residual, it cannot be used for prediction on the test data. Panel **(D)** shows HA residuals and response predictions of the tide-surge response for different values of tide and surge. Note, Panel (D) is truncated between surge values of 1 and 2.2 meters for legibility. RTide uses surge values as an additional input function.

least squares (IRLS). While IRLS improves performance over OLS for contaminated signals, the RTide tidal prediction still yields a reduction in MAE from 0.010 m to 0.009 m.

Due to the fact that our system is separable, one may assume that the linear inverse barometer effect could be estimated directly by simply regressing the atmospheric pressure $P$ against the non-tidal residual (surge + tide-surge interaction). This procedure assumes the remaining variability after removing the harmonic or response computed tidal signal is Gaussian distributed noise. Figure 5.2 Panel(s) B compare the HA and RTide estimated inverse barometer effect (IBE) by simple linear regression of the non-tidal residual with the atmospheric pressure. The impact of the biasing in Panel A is clearly shown. While RTide recovers the exact IBE of -1.002

cm / mbar, the HA obtains an IBE of -0.975 cm/mbar. While 2.5% error may seem minor, these errors compound with the tidal errors and any residual noise in the observations when estimating the tide-surge interaction.

Figure 5.2 Panel(s) (C) shows how the biasing of the derived HA tidal signal significantly degrades the accuracy of the residual tide-surge interaction. Here, we are comparing the learned and predicted tide-surge interaction from RTide with the tide-surge interaction obtained by subtraction of the HA tidal prediction. It can be seen that the RTide learned and predicted tide-surge interaction converges to the true system, as is evidenced by $R^2$s of 0.999 and 0.923 on the train and test sets respectively. The degradation in the RTide estimated tide-surge interaction between train and test-set is largely from errors at low tide-surge values in which tidal errors also contribute. In comparison, the HA residual tide-surge interaction shows an $R^2$ of 0.883 for the train set indicating that the biasing of the tidal prediction results in a degradation of the true nonlinear system once subtracted. Since this is only a residual and not a prediction, an equivalent test set prediction is not possible. The degradation of the HA $R^2$ is primarily due to errors at low surge values but is also a function of the noise of the data. Given the low amount of Gaussian noise applied, this represents a best-case scenario. As shown, even infrequent non-stationary contamination can yield systematic tidal estimation error. This is further confirmed by Figure 5.2 Panel (D) in which the RTide demonstrates excellent agreement with the actual nonlinear system and the HA is severely biased. These results indicate that RTide can learn the true nonlinear impulse-response relationships directly from the data. This result is important for application on real-data where studying these relationships using conventional parametric methods is challenging. Furthermore, we have demonstrated the shortcoming of using standard harmonic methods to study systems contaminated by non tidal forcing. We stress that any analysis conducted on the biased residual will result in similarly biased conclusions.

Similar results are observed when comparing the accuracy of the derived harmonic constituents. The error of the $k^{th}$ constituent in terms of the quadrature amplitudes is estimated using the RMS (23), defined as

$$\mathrm{RMS}_k = \sqrt{[(A_{\mathrm{model}}\sin(\omega_k t) + B_{\mathrm{model}}\cos(\omega_k t)) - (A_{\mathrm{true}}\sin(\omega_k t) + B_{\mathrm{true}}\cos(\omega_k t))]^2}. \tag{5.8}$$

This metric provides a balanced measure of the derived harmonic constituent error by penalizing both amplitude and phase proportionally. Results are provided in Table C.2 in Appendix C. By including the atmospheric pressure component as an additional input into RTide, a lower RMS error of the tidal constituents is obtained for a majority of constituents. An estimate of the total error across all constituents is given by the root-sum-of-squares (RSS), where RSS = $\sqrt{\sum_{k=1}^{n}(\mathrm{RMS}_k)^2}$. The RTide derived tidal signal shows a lower RSS over both OLS and IRLS harmonic analysis methods. While the improvements over IRLS are only slight for this example, it is expected that these disparities will grow in the presence of a continuous noise spectrum. Regardless, we emphasize that the optimal choice of method depends on the objectives of the researcher and that an optimal analysis may result from using both RTide and HA. Section 5.4.2 shows how RTide can be used to learn the nonlinear response to meteorological forcing for the more challenging problem when the surge is generated by multiple meteorological processes.

## 5.4 Real Data

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

### 5.4.1 Tidal Rivers

Sitting at the boundary between hydrology and physical oceanography, tidal rivers represent a unique challenge in tidal modeling. The confluence of river outflow,

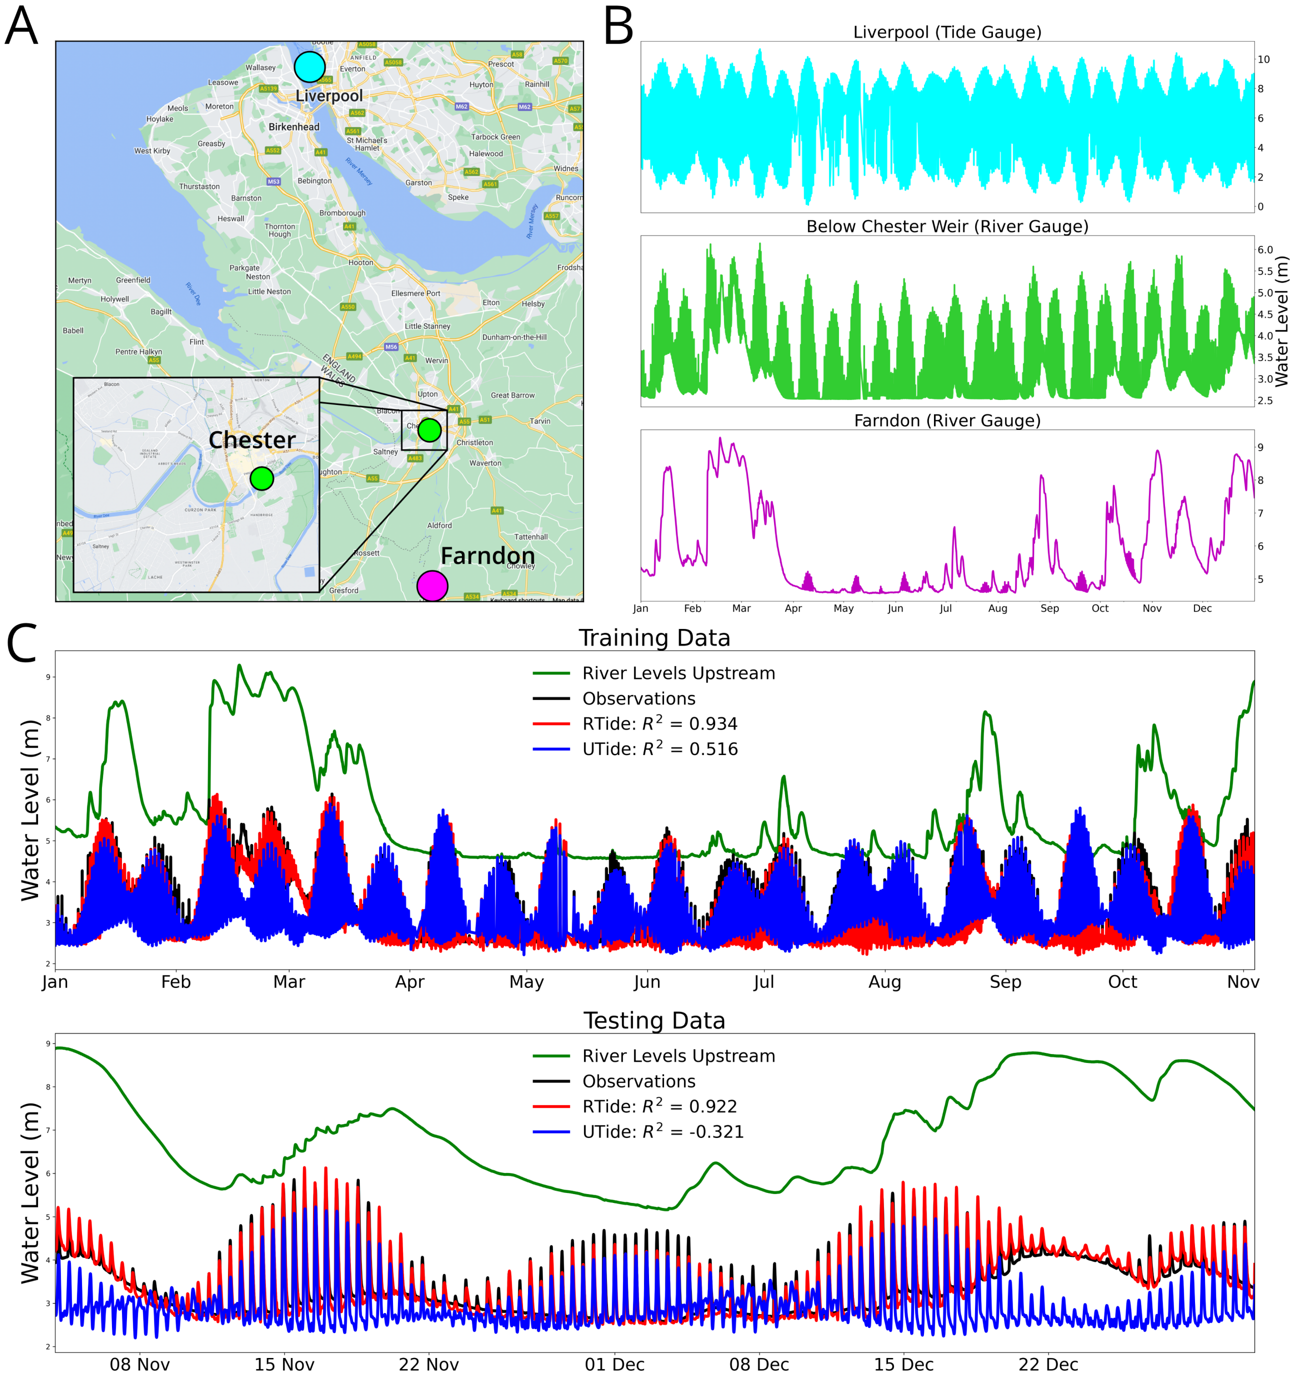

Figure 5.3: **Comparisons between RTide and IRLS HA predictions of tidal river levels beneath Chester Weir, UK.** Panel **A** shows the relative location of each data source. Panel **B** shows the time series of each data source. Panel(s) **C** show RTide and UTide predictions over the full year with the top panel being hindcasts (Training Data), and the bottom panel showing forecasts (Testing Data). The coefficient of determination, $R^2$, of each models predictions is given. The only inputs into RTide are the upstream river levels, and gravitational input functions.

variable bathymetry, and sediment input reduce the performance of conventional harmonic and numerical methods (150). Despite this, modeling these regions has significant implications for the global water budget, and in regional resource management and flood mitigation efforts (151). We demonstrate the improved ability of RTide to capture this behavior and to integrate multivariate data sources using data from beneath the Chester Weir, near Liverpool, UK for the year 2020. As a tidal river, the observed water levels at this location are a consequence of the nonlinear interaction between the incoming tide and the outflow of the River Dee.

While non-stationary HA methods have been developed for the prediction of tidal rivers, they require the presence of a secondary reference station free from non-stationary contamination which significantly reduces their global applicability (35; 118; 152). Additionally, the basis functions employed in NS_TIDE (118) require river outflow measurements, and not river water-levels. For many locations in the UK and globally, river outflow measurements are unavailable. To combat this, we use upstream water levels as direct input into our model, and use RTide to approximate the localized sea-level response. While river levels are not as accurate as river outflow, they are much cheaper to obtain and can even be extracted via satellite data. As such, the proposed approach removes the current requirement of outflow data and can be applied globally. It is in theory possible to define a basis function for NS_TIDE that could take upstream river levels as input. However, to the best of our knowledge, this has not been done in the literature and would require significant expertise input. Future work will look to analyze the learned RTide model to produce an approximate basis function that can be integrated into the NS_Tide framework for application to river level data. We note that this would not remove the need for a secondary port in NS_Tide which we regard as a significant innovation of our approach.

Figure 5.3, Panel (A) shows the relative location of each of these data sources. Figure 5.3, Panel (B) shows the observed sea level below the Chester Weir, alongside concurrent tide gauge data from Liverpool, and measured water levels of the River Dee upstream of the weir. While there is a relationship between the upstream river level, tidal signal, and the water elevation below the Weir, the relationship is not linear and crucially, non-harmonic (118).

To test the performance of both methods, RTide, and IRLS HA are trained using 10 months of reference data (≈20,000 observations) from beneath the Chester Weir. Here, IRLS HA is utilized as the data is strongly non-stationary. Predictions are then computed for the subsequent 2 months (November to December). The 10 month hindcasts and 2 month forecasts are shown in Figure 5.3, Panel (C). In addition to the spherical harmonic input functions, RTide is provided with the real-time upstream river levels. Critically, unlike other nonstationary methods, RTide does not require the inclusion of an adjacent gauge without fluvial effects.

Visually it can be seen that the RTide prediction is better able to track the non-stationary fluctuations induced by river outflow than HA. This is confirmed by the coefficients of determination ($R^2$) of 0.922 and −0.321 for RTide and HA, respectively, on the test dataset. This result is not surprising as the tidal dynamics are heavily modulated by the tidal river as will be shown in the subsequent SHAP analysis. By removing the assumption of stationary behavior, RTide can easily learn the coupled response of sea-levels to non-stationary forcing. We note that the use of IRLS HA in this case study is meant to illustrate the severe biasing of harmonic predictions when applied to areas exhibiting nonstationarity and is not meant as a direct benchmark.

RTide demonstrates balanced performance across the training and test sets with $R^2$ values of 0.934 and 0.922 respectively (the difference is not significant based on t-test with $\alpha = 0.05$). This indicates that RTide has learned the localized and nonlinear response of the sea-level beneath the Chester weir to the combined tidal and fluvial forcing. The learned model can therefore be used to inform decisions and quantify the risk associated with different water management scenarios. This is important as the gauge beneath the Chester weir has since been decommissioned. Unlike autoregressive machine learning approaches (153; 128), our approach enables the accurate and continued prediction of water-levels in this region based on this historical data.

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

This can be seen in Figure 5.4, and Figure C.3 in Appendix C. Here the RTide model is analyzed using the SHAP analysis procedure outlined above. Dependence plots are shown which visualize the learned response *distributions* between the individual input functions and the associated sea-level response beneath the Chester Weir. These distributions are produced by interactions between input functions. Indeed, if no such interactions exist then an analytic form could be approximated, but care should be taken.

Each plot illustrates the individual contribution of each physical driver, and the nonlinear interaction between them. As noted, the spread in the derived response relationships is a consequence of interactions with other inputs, which when plotted together, can shed light on how these physical drivers interact. For example, Figure 5.4 shows the learned response to the upstream river levels, and exhibits a near cubic relationship ($R^2 = 0.987$) with relatively little spread in the associated distribution reflected by the ratio of the standard deviation of the fit model's residuals $\sigma_{fit}$ compared to the standard deviation of the data $\sigma_{data}$ ($\sigma_{fit}/\sigma_{data} = 0.11$). Inspection of the distribution of the color-mapped astronomical tide indicates that at low upstream river levels ($\leq 7$ m) interactions with the tide create a larger sea-level response. In contrast, for high upstream river levels ($\geq 7$ m), and thus high river outflow, the incoming tide appears to have a damping effect. These learned responses nicely align with theory (129).

Figure 5.4 illustrates the impact the river outflow has on the astronomical tide by comparing the RTide derived gravitational response between Liverpool and Chester for individual input functions and time-lags. $\text{Gravitational}_n^m$ refers to the $m^{th}$ order and $n^{th}$ degree gravitational spherical harmonic input function. The $\text{Gravitational}_2^2$ spherical harmonic shown in Figure 5.4 corresponds to the semi-diurnal species and thus primarily drives its response. The tides at Liverpool have some fluvial influence due to the river Mersey but sit adjacent to the open ocean and can thus

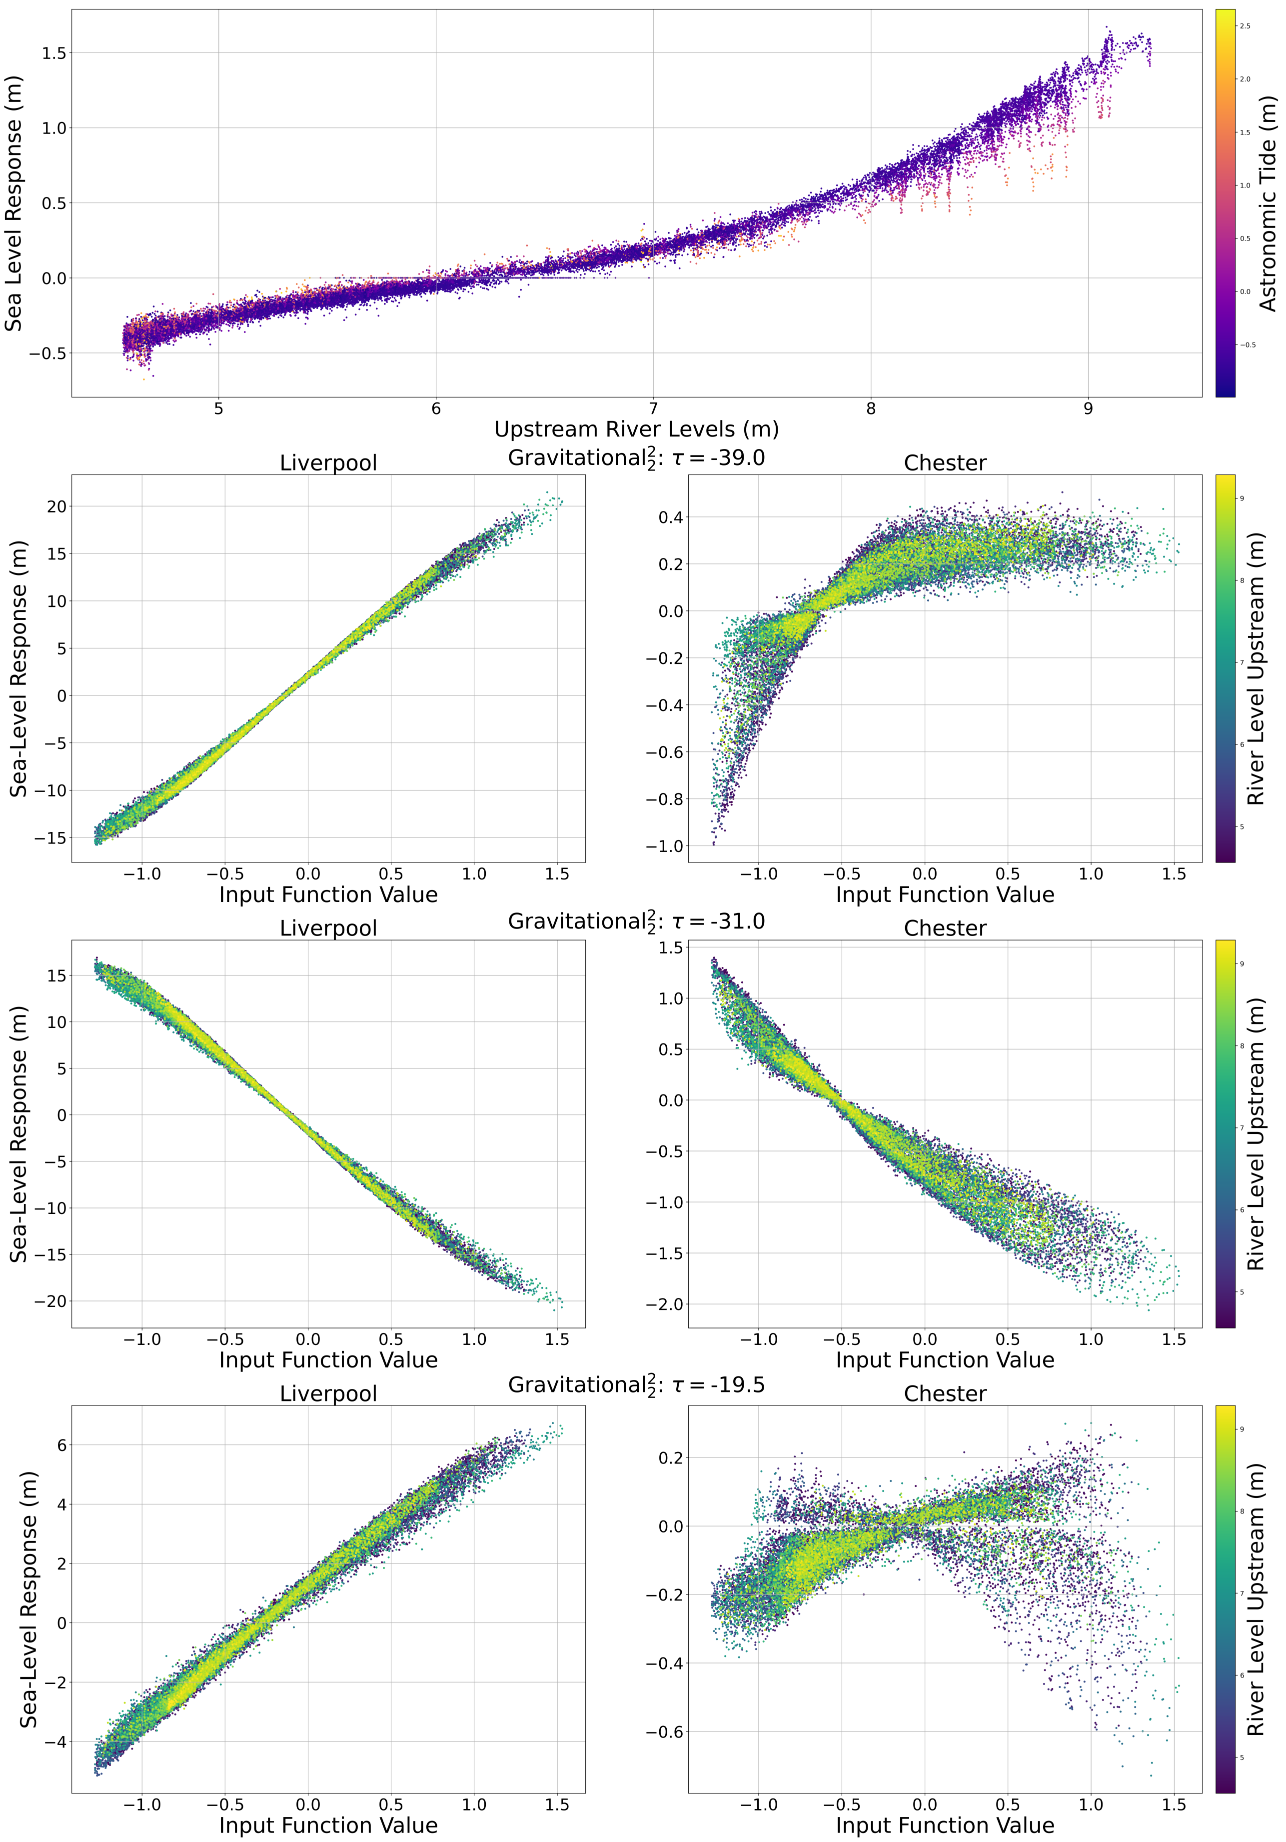

Figure 5.4: **SHAP derived response distributions of the Chester, UK sea-level. (Top Panel)** Learned response distribution to upstream river levels. Data points are colored according to the model predicted astronomical tide. **(Bottom Panels)** Comparisons of learned response distributions to the second order and degree Gravitational input functions at associated time-lags $\tau$ between Liverpool, UK and Chester, UK. Chester data is colored according to the upstream river levels at the time of measurement. Corresponding input function, spherical

be used to evaluate the impact of the introduction of river outflow. It can be seen that the river outflow modulates the tidal response in several ways. First, even for time-lags associated with the semi-diurnal tidal response, we see the upstream river levels induce a spread of the sea-level response for the Chester gravitational response distributions, as opposed to the relatively minimal spread of the Liverpool responses. This spread exhibits a clear gradient when considered with the river outflow. Here, high river levels upstream appear to damp the gravitational response leading to larger sea-level responses at lower river levels and vice versa.

The second factor affecting the response distributions is nonlinear distortion. The Liverpool responses can be seen to be near linear based on linear regressions with $R^2 > 0.99$. In contrast, the Chester response distributions are distorted and nonlinear (regression $R^2$ of 0.62,0.93, and 0.16). Figure C.3 in Appendix C shows the effect of these modulations on the overall gravitational input functions separated by different ranges of upstream river level. These varied nonlinear relationships highlight the difficulty encountered in the classic approach – where one must prescribe these relationships *a priori*. Future work will look to explore these relationships further, and to connect the observed response distribution modulation to the extensive work which has been done for HA (129; 118).

An interesting aspect of the gravitational sea-level response distributions in Figure 5.4 is the location of the zero-point of the sea-level response or "waist". The distributions for both Liverpool and Chester appear to narrow about this point. Interestingly, at Liverpool the center point appears to occur close to an input function value of zero regardless of time-lag. In contrast, the waist of the Chester distributions for lags $\tau = -39$ and $\tau = -31$ appears to shift to the left (occurring closer to -0.5). Several possible explanations exist which warrant further investigation. First, it is possible that the shift in the waist is simply related to the difference between the timing of the tides between Chester and Liverpool. This could effectively be ruled out using several stations in succession. An alternate explanation is that the decoupling of the waist occurs due to frictional effects (150).

While a river-level forecasting model is needed to operationalize these tools, consid-

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

Table 5.1: Comparisons of HA and RTide performance from January 1, 2015, to December 31, 2018 in Money Point, VA (NOAA Station ID: 8639348). Here "Tidal" refers to RTide using just gravitational and radiational input functions. Additional models are named according to the additional meteorological variables they are trained on. Best values are provided in bold. Statistically significant improvements in $R^2$ over HA are indicated using * and are computed via t-test with $\alpha = 0.05$.

erable work in the field of hydrology has looked to solve this problem (154; 155; 156). Furthermore, while results could be improved by using river outflow as an additional forcing input, the proposed approach has the advantage of being applicable to remote-sensing data. This will be useful for the development of global estuarine and tidal river models on the upcoming NASA Surface Water Ocean Topography (SWOT) mission. Further study is needed using data where both river outflow measurements and uncontaminated secondary ports are available to directly compare the predictive capabilities of RTide with NS_TIDE.

### 5.4.2 Meteorological Forcing

As sea levels continue to rise, coastal flooding and storm surges will become more frequent (157). While major flooding events, primarily driven by hurricanes and typhoons, are of great importance, the incidence of minor flooding will proliferate in the coming years. Often referred to as nuisance flooding, these events rarely result in loss of life but place severe strain on public infrastructure and property (158). In fact, the cumulative cost of these frequent events can exceed that of less frequent, major flooding events (159). The drivers of nuisance floods are generally more subtle and are a result of the confluence of astronomical tides, wind, and weather.

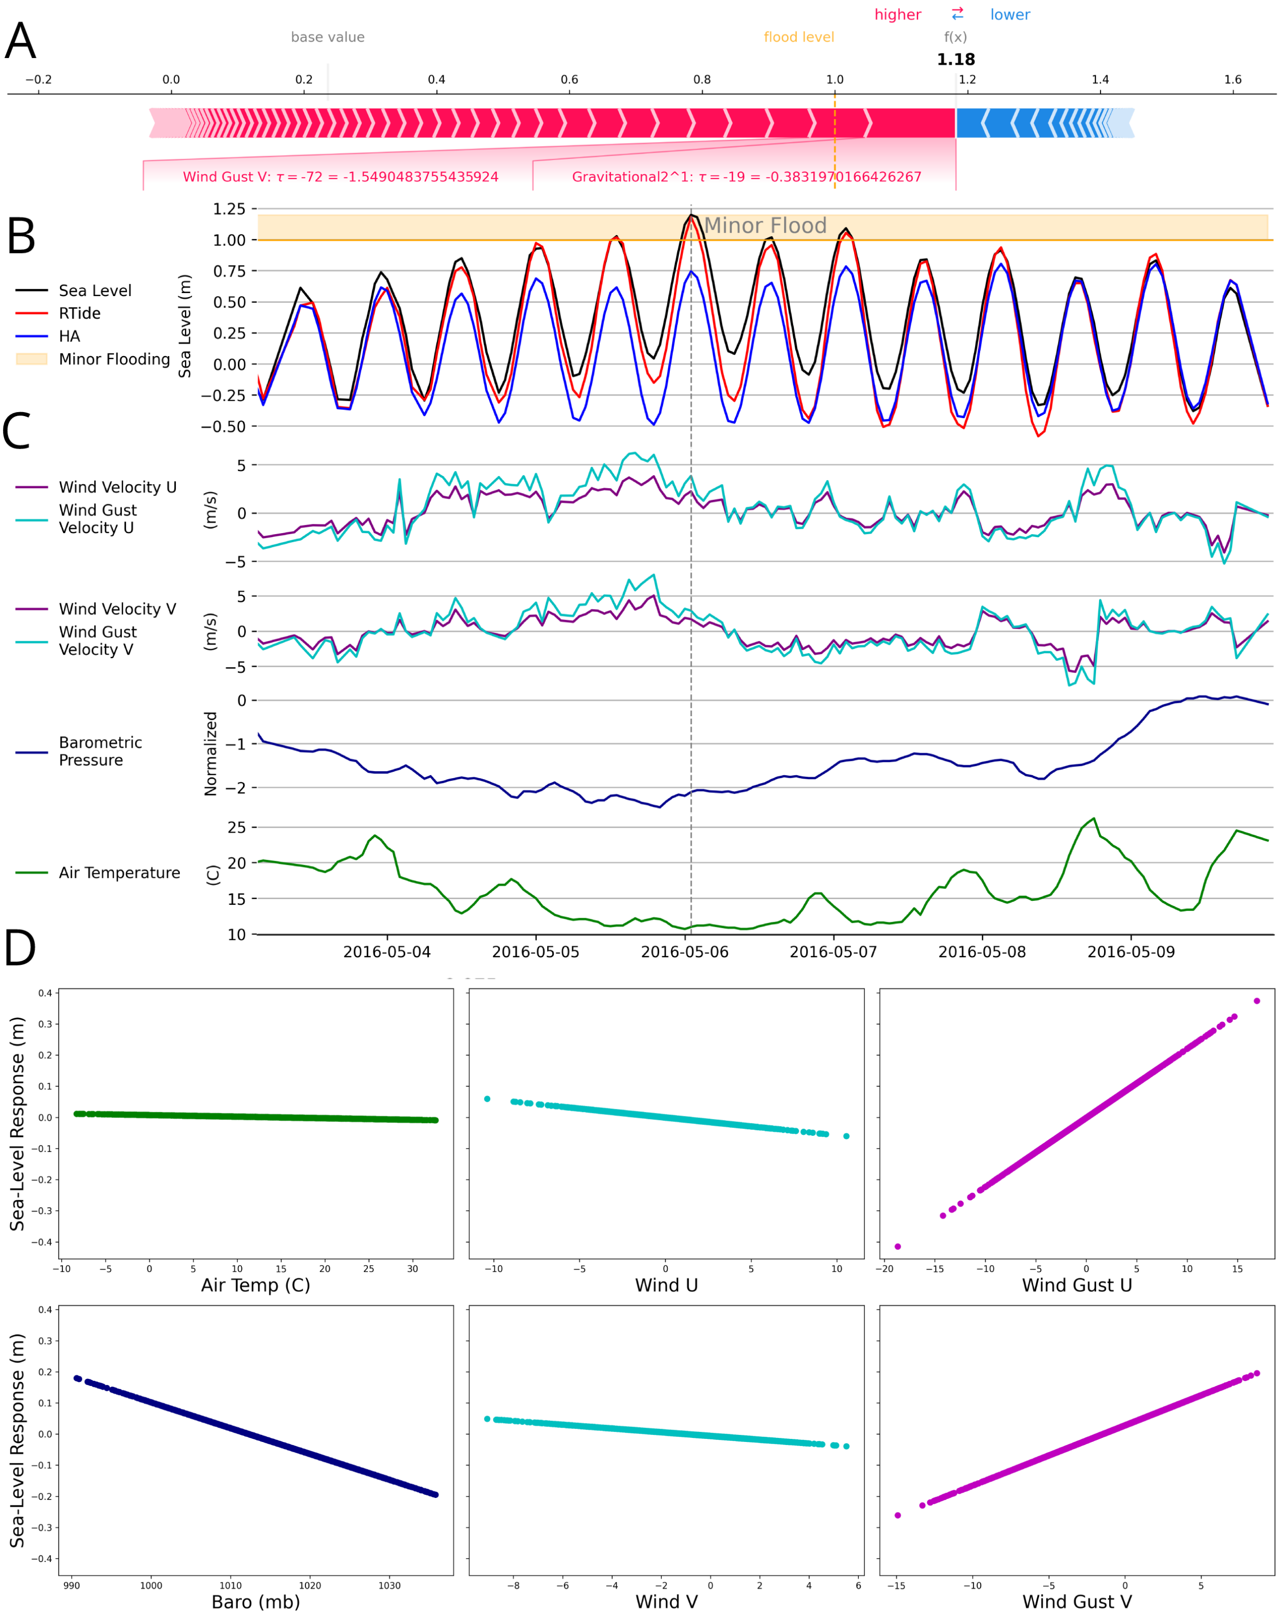

Figure 5.5: **5-day nuisance flooding event in Money Point, VA, USA.** Panel **A** shows the explained surge prediction at the time instance labelled 'Minor Flood'. Panel **B** shows RTide and HA predictions of the observed sea level. Panel **C** shows the subset of meteorological variables associated with the 5-day flooding event and RTide predictions. Panel **D** shows distributions of SHAP sea-level dependence by real-time meteorological variables. Distributions are shown for 4200 predictions over 180 days from January through June of 2016. RTide uses all meteorological variables shown.

Classically, the task of surge prediction is carried out using numerical methods (160). These hydrodynamic models are computationally intensive and suffer from over-simplistic assumptions that prevent the consideration of coupled processes (161; 162). In particular, many operational models treat tides and surges as separable due to computational limitations. Several studies have shown that treating surge and tide as separate can lead to erroneous predictions (163; 164; 165). Additionally, the high computational cost of numerical models is impractical for ensemble forecasting and regions with poor data quality or access to computational resources. To combat these limitations we propose the ML Response Framework as a lightweight tool for learning the localized response to meteorological forcing. We stress that the present study is not developing an alternative to the operational surge models, but rather illustrating the potential of our approach for further applications in surge modeling and understanding the role of meteorological forcing.

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

Evaluation of the SHAP values for the model's prediction of a flooding event can help to explain the event's physical drivers. We illustrate this using a 5-day surge event from May 4th to 9th, 2016, characterized by sustained on-shore winds. The minor flood level threshold is designated by NOAA's Weather Forecasting Office and is empirically defined by occurrences of local flooding. For Money Point, VA this is 54 cm above the mean higher high water. Figure 5.5, Panel (A) explains the influence of each forcing input on the RTide prediction of the first minor flooding event when trained on pressure, wind, temperature, and wind gusts. It can be seen that the combination of high astronomical forcing and sustained southerly winds over the previous 72 hours contributed to this event. This information can be useful to local environmental agencies as the tide-surge response to meteorological forcing is highly location and time-dependent.

Panel (B) compares RTide and HA predictions for the full 5-day surge event. Visually, the RTide prediction shows strong agreement with the observed sea level, particularly during the peak flood events. While predicting these occurrences is clearly of the highest importance, prediction accuracy decreases during low-tide events. This is likely a consequence of only considering localized meteorological forcing. As was demonstrated in Section 5.4.1, we anticipate the inclusion of additional spatial information will further improve predictive accuracy. Experiments with the inclusion of positive meteorological lags showed promise in predicting external surges (e.g. surges originating from elsewhere on the shelf). However, further analysis is needed to confirm these findings.

Panel (C) shows the evolution of the four meteorological variables over the 5-day surge event horizon. These reveal a key challenge in predicting nuisance floods; flooding results from the coincidence of high tide events with both the chronic build-up and acute effects of meteorological forcing. The acute impacts of meteorological forcing can be seen in Panel (D), which shows distributions of the model's learned dependence of the sea-level to real-time meteorological forcing taken over a 6 month period. Analysis of these distributions shows that the learned RTide feature importance is consistent with our physical intuition. There is a clear, and significant, inverse linear relationship between barometric pressure and the localized sea-level response. Linear regression of the learned sea-level response to pressure yields a slope of $-0.83$ cm/mbar which is in good agreement with both theoretical and empirical studies (149). Some discrepancy may arise from the fact that we only consider the barometric pressure at the gauge, and not the average over the region. Additionally, we see that elevated wind gusts in the south-westerly direction have the most significant impact on real-time sea-level. This finding makes sense as the station sits on the South-Western point of a jetty. It is interesting that the wind exhibits slight negative sea-level responses (-.6 cm/kmh$^{-1}$), and the gusts large positive responses (2 cm/kmh$^{-1}$). However, a simple explanation is that the total response to wind is the sum of the response from both sustained winds and gusts. Adding these two components shows a total wind response which is approximately linear with a response slope of (1.14 cm/kmh$^{-1}$) for the U component (Figure C.4 Appendix C). The spread in the distribution (nonlinearity) seems to increase with the difference between sustained wind and gust as shown by the colormap. The model response shows ambient temperature has essentially no impact ($-0.05$ cm/C$^\circ$ ). This illustrates the value of SHAP analysis in explaining models with multiple partially correlated inputs which proved difficult to define or interpret in Cartwright's original surge analysis (16). Through our testing, we found these relationships became more clearly defined when models were trained on longer reference series and regularized. Future work will look to verify the nature of these dependencies on longer reference series and to compare them to different stations.

In this application, we use the observed meteorological variables as direct inputs into our model. For an operational model, regional forecasts would need to be substituted. This ‘lightweight’ surge-prediction framework shows promise as an alternative to conventional numerical methods for communities lacking computational resources and in predicting nuisance flooding. Notably, none of the examples presented here utilize the observed water-levels as input in contrast with the machine learning approaches previously employed for this task(128). Inclusion of recent or real-time water levels as exogenous inputs, termed ‘self-prediction’ by Cartwright (16), can significantly improve performance. However, by avoiding this requirement, our approach is readily applied to ungauged locations and does not suffer from potential observational issues.(128) Additionally, while such an approach naturally falls out of the response framework, it is not possible in a harmonic context.

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

## 5.5 Discussion

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

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

For "messy" applications in more complex tidal regimes, the flexibility and non-parametric formulation of our method become advantageous. As illustrated in Section 5.3.1, conventional HA can cause problems when applied to signals contaminated with non-astronomical forcing. One such example is the biasing of harmonic coefficients from radiational effects. Known as the radiational double-counting problem, this can lead to significant errors for operational surge forecasts (119; 166). While non-stationary extensions to HA have been developed and applied for the analysis and prediction of sea-levels in tidal rivers and estuaries (118; 35), these approaches encounter several challenges. First, both of these approaches require an adjacent 'uncontaminated' reference series. As noted above, for many regions globally, particularly in the Global South, this is not feasible. Additionally, with regard to NS_Tide one must first understand the physics of the interaction between tide and external forcing, which, for many important problems, is not yet well defined (storm surges, sea level rise, intertidal regions, etc.).

The non-parametric nature of the ML approach removes this requirement and enables the study of these phenomena without imposing any assumptions regarding separability, stationarity, or the structure of interaction kernels. The flexibility in input forcing selection garnered by these characteristics makes the testing of multivariate forcing simple and efficient. Future work will look to combine these capabilities with symbolic regression in order to derive analytic approximations of the learned response models. The insights could then be infused into conventional non-stationary HA to leverage the aforementioned benefits of a simple regression model. We remark that an analogous framework could be conceived of for HA if a non-parametric term could be simultaneously estimated (e.g. NS_Tide without predefined basis function), similar to the use of 'universal differential equations' (167). As shown in Section 5.3.1 we believe that simultaneous estimation is critical to recovering the true physical processes.

In addition to the results shown in Sections 5.4.1 and 5.4.2, the approach to analyzing the learned model using SHAP provides many further avenues with which to study the impacts of non-gravitational forcing. Indeed, Figures 5.4 and C.2 in Appendix C provide a first look at how this can be used to understand how non-gravitational forcing modulates the astronomical tide and *vice versa*. Additionally, Section 5.4.2 demonstrated how the SHAP values can be used to recover known physical processes, such as the inverse barometer effect, even in the presence of multiple drivers. This constitutes an exciting new area of study as countless phenomena interact and modulate the astronomical tide but are difficult to study and predict using conventional methods. For example, interactions of tides with mean sea-level (and thus implications for sea level rise) (16), and even the impact of barrages, are both highly salient for coastal resilience efforts. However, current efforts to include these interactions in engineering practice are generally limited to statistical estimation or computationally intensive numerical simulation (168). Initial testing has shown our approach to be promising for both such applications and will be pursued in the future. Given their physical interpretability, we believe SHAP values can help enable new explorations into the dynamics of the ocean's response to various forcing, as demonstrated.

The present study has primarily conducted our analysis of the learned RTide models within the time-domain. As shown in (5; 16), the frequency interpretation of response models is similarly rich. In fact, Munk and Cartwright found the admittance function $Z(f)$ of the learned response weights to be more meaningful than the response weights themselves. It is possible to obtain the equivalent Volterra series directly from the trained networks (133). This approach could be useful in

(i) studying the nature of the learned interactions and (ii) translating these models back into a familiar response regression. The learned admittance $Z(f)$ can then be computed and analyzed in the standard way. Alternatively, as described in Section 5.2.3, SHAP values can serve as a transformation from our learned ML model to a generalized linear model of binary variables. Hence, future work will look to exploit both of these approaches to analyze the approximate admittance of learned RTide models and compare them with the original method.

The accurate prediction of estuarine tides remains a significant challenge due to distortion of the tidal signal which can give rise to asymmetric and non-harmonic profiles (sawtooth, bore formation, etc.). Inaccuracies from these predictions constitute one of the largest sources of error for the UK operational surge model. (166). As demonstrated, a response approach can readily capture and predict these modulations. Future work will look to operationalize these forecasts for integration into these tools.

## 5.6 Conclusion

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

By utilizing a class of neural networks which are equivalent to the Volterra series, the approach retains equivalence to the classical response method, and under certain conditions, the analytic form can be recovered exactly (135). Constraining these networks within the impulse-response framework:

1. Removes the need to define interactions between input functions.
2. Allows the network to learn the higher order impulse-response relationships.
3. Enables prediction of non-stationary sea-levels without requiring an adjacent reference station removed from such influences.

4. Makes easy the inclusion, quantification, and analysis of interactions between non-astronomical forces with tides.

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

## Open Research Section

NOAA tide gauge data (section 5.4.2) can be accessed freely via NOAA Tides and Currents repository tidesandcurrents.noaa.gov/. All data and code needed to reproduce the figures are provided in the given respository. The tide and river gauge data in section 5.4.1 is provided courtesy of the UK National Oceanography Centre in the provided respository. Code to generate the synthetic surge data is provided in the RTide repository.

**Code Availability:** RTide (`github.com/thomasmonahan/RTide`) is a public Github repository and Python package (169). It can be installed via the commandline using pip, or through cloning the respository directly. Jupyter notebooks containing the model implementation and subsequent analysis are included for all figures.

# 6 Response-based prediction of tidal currents

# 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. This performance advantage increases with the complexity of the tidal currents. More modest reductions in absolute error are observed for less-complex sites with the response approach reducing predictive error by 9.6% on average across 40 active NOAA current stations. The framework offers several opportunities for future work in understanding the role of non-tidal forcing and sediment transport and has significant economic implications for tidal energy site development. The proposed approach is implemented and freely available in the open-source RTide Python package. Critically, the non-parametric approach reduces the need for expertise when applying the response method to study higher-order nonlinear processes.

## 6.1 Introduction

Tidal currents present unique challenges for tidal analysis and prediction (10). 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 (170; 171; 20). Unlike tidal elevations, tidal currents exhibit significant vertical and horizontal variability (172). 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 (173).

Classically, analysis and prediction of tidal currents are conducted with harmonic analysis (HA) (10). 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

$$\hat{\zeta} = \sum^{N} A_k \sin \omega_k t + B_k \cos \omega_k t, \tag{6.1}$$

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 (33; 34; 105). 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 (105). Typically, constituent selection is done using the Rayleigh Criterion which defines the "separability" of two constituents with frequencies $\omega_1$ and $\omega_2$ based on the time-series duration $T$ such that

$$\left(\frac{T}{\omega_1 - \omega_2}\right)/R_{min} \geq 1, \tag{6.2}$$

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 (1). While more advanced constituent selection procedures have been developed (45), 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 (174). This limits the number of harmonic tidal constituents which can be estimated and thus the accuracy of the predictions (1). 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 (20).

With these operational challenges in mind, it is worth remarking from the outset what this chapter will address. This chapter 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. There have been extensive contributions to the tidal currents literature which devise methods for analyzing these signals and separating out nonstationary components (44; 15; 45; 130). 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 (175; 176). 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.

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 (172; 173; 10). 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 (17). Unlike the classical response method (5), 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 6.2.7).

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.

## 6.2 Methods

### 6.2.1 Response Method

The response method exploits the fact that the oceanic response to forcing is weakly nonlinear and time-invariant (5). 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 viewed 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 6.5. 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

$$V(\theta, \lambda; t) = g \sum_{n=0}^{n} \sum_{m=0}^{n} \left[ a_n^m(t) G_n^m(\theta, \lambda) + b_n^m(t) H_n^m(\theta, \lambda) \right], \tag{6.3}$$

where $g$ is the gravitational constant, $\theta$ is the Geographical Colatitude, $\lambda$ 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 $V(\theta, \lambda; t) \in C$. The complex spherical harmonics $G_n^m + iH_n^m$ converge rapidly and are defined as

$$G_n^m + iH_n^m = (-1)^m \left[ \frac{2n+1}{4\pi} \right]^{\frac{1}{2}} \left[ \frac{(n-m)!}{(n+m)!} \right]^{\frac{1}{2}} P_n^m(\cos\theta) e^{im\lambda}, \tag{6.4}$$

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 (5), was found useful to deal with seasonal effects, we find it degrades performance when performing long-term predictions and it is thus not considered further.

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

$c_n^m = a_n^m + ib_n^m$ such that

$$\hat{\zeta}(t) = \sum_{m,n} \sum_{s} [x_n^m(s) a_n^m(t - \tau_s) + y_n^m(s) b_n^m(t - \tau_s)]. \tag{6.5}$$

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

$$i^{th}\text{order response} = \sum_{i} \cdots \sum_{x} \sum_{s} \cdots \sum_{s'} w(i, \ldots, x, s, \ldots, s')(c(t-\tau_s))(\ldots)(\ldots)(c(t-\tau_{s'})). \tag{6.6}$$

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 6.5, 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. Here, species 1 and 2 correspond to the diurnal and semi-diurnal response, respectively. 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 (177). As noted in (17), 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 6.2.4.

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 (5; 16; 177; 39). Second, nodal corrections are implicit (e.g. embedded in the forcing itself), which greatly simplifies long-term tidal analysis and prediction tasks. This allows the learned response function to explicitly model departures from equilibrium nodal theory which is necessary in many regions globally (178). 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 6.2.7. 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.

The response weights themselves are not as interpretable as the amplitudes and phases of HA (5). As described below, the admittance function (Fourier transform of the weights) in a response analysis is generally of greater interest.

### 6.2.2 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.

The coupled response model for two orthogonal components of tidal currents can be written as

$$\hat{\beta}(t) = \begin{bmatrix} \hat{u}(t) \\ \hat{v}(t) \end{bmatrix} = \sum_{m,n} \sum_{s} W_n^m(s)\, c_n^m(t - \tau_s), \tag{6.7}$$

where the weights $W_n^m$ correspond to the matrix

$$W_n^m(s) = \begin{bmatrix} x_{n,\vec{x}}^m(s) + y_{n,\vec{x}}^m(s) & x_{n,\vec{x}\to\vec{y}}^m(s) + y_{n,\vec{x}\to\vec{y}}^m(s) \\ x_{n,\vec{y}\to\vec{x}}^m(s) + y_{n,\vec{y}\to\vec{x}}^m(s) & x_{n,\vec{y}}^m(s) + y_{n,\vec{y}}^m(s) \end{bmatrix}, \tag{6.8}$$

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 6.3.3 provides further evidence for this innovation and considers the improvements it garners in predictive accuracy.

### 6.2.3 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 (5).

As in (17), 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. 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 (5).

### 6.2.4 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. Here, the frequency spectrum's $H$ and $G$ are computed for the observed sea-level $\zeta^j$ and a given input function $\gamma^j$ respectively such that

$$H^j(f) = \int_{-\infty}^{\infty} \zeta^j(t) e^{2\pi i f t} \mathrm{d}t, \tag{6.9}$$

and

$$G^j(f) = \int_{-\infty}^{\infty} \gamma^j(t) e^{2\pi i f t} \mathrm{d}t. \tag{6.10}$$

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

$$\hat{\zeta}^j(t) = \int_0^{\infty} \gamma^j(t - \tau_s) w_n^j(s) \mathrm{d}\tau. \tag{6.11}$$

In the frequency domain, this is equivalent to

$$H^j(f) = Z^j(f) G^j(f) \tag{6.12}$$

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

$$Z^j(f) = \int_0^{\infty} w_n^j(\tau) \exp(2\pi i f \tau) \mathrm{d}\tau. \tag{6.13}$$

This is the tidal *admittance* function (179). 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 (5; 16). 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.

One can also perform a least-squares type procedure as in modern HA. Here, the response prediction $\hat{\zeta}$ can be cast as a linear model with $\hat{\zeta} = w x_i + \epsilon$. 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 (17) 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 (133; 135). 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 $\tau$ 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 6.7, can be replaced by

$$\hat{\beta}(t) = \begin{bmatrix} \hat{u}(t) \\ \hat{v}(t) \end{bmatrix} = f(V_0(t-\tau_0), \ldots, V_0(t-\tau_S), \ldots, V_X(t-\tau_0), \ldots, V_X(t-\tau_S)), \quad (6.14)$$

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 6.2.7.

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 as only learning interactions for which there is sufficient evidence in the data to support. 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.

### 6.2.5 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):

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

This can be seen to place a greater weighting on outliers which can lead to the biasing of the learned model. Similar to what (34; 33) 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 $\delta$.

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

By default, we utilize $\delta = 1$. This was found to improve performance for the operational applications in Section 6.4 which exhibit significant noise. Since the data is normalized during training, we find this choice of $\delta$ 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.

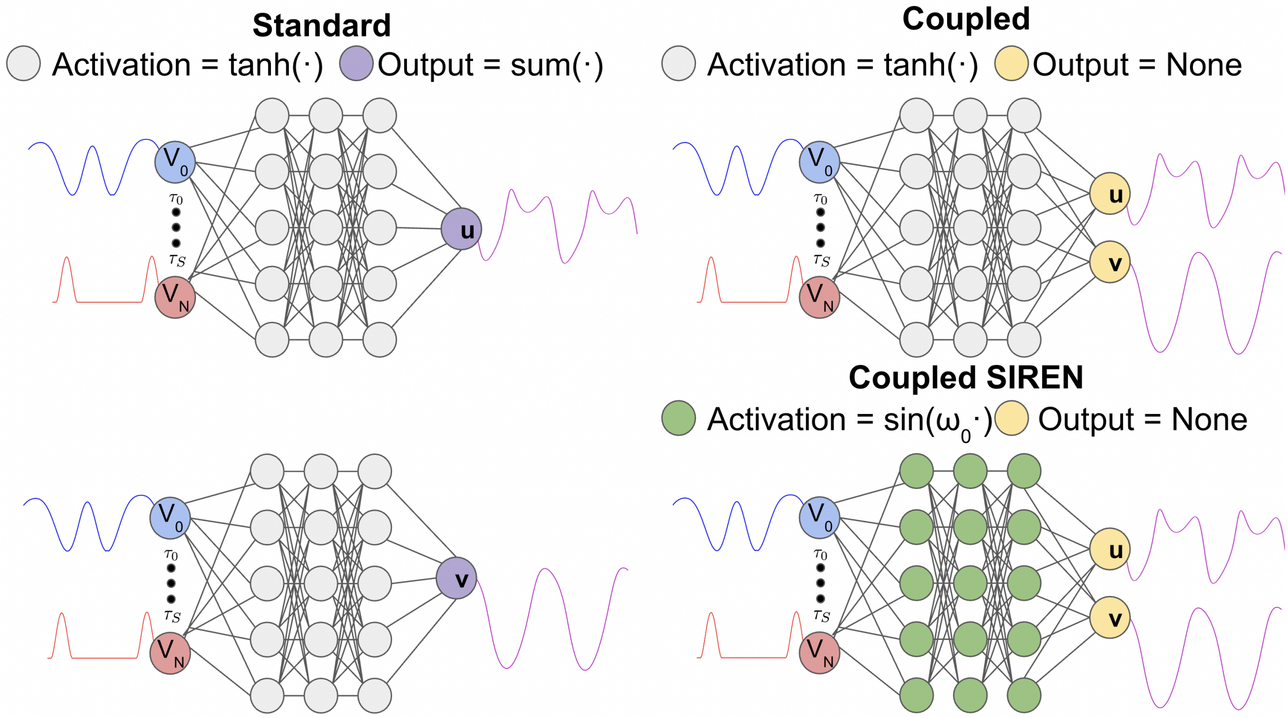

Figure 6.1: **Comparison of RTide model architectures.** Activation refers to the activation function used in the hidden layers of the network. Output refers to the activation function used for the ouput node(s) of the network. If the activation = None, no nonlinearity is applied and thus the final output is simply given by $wx + b$ where $x$ is the input and $w$ and $b$ are the learned weight and bias.

### 6.2.6 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 and nodes, 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. 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 6.1.

#### 6.2.6.1 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 (17). The given architecture has 3 hidden layers with the number of neurons in each layer equal to the number of inputs. The final output consists of a simple "Adder" layer which sums all outputs. Tanh activation functions are employed to yield equivalence to the Volterra series (135). By treating the $\vec{u}$ and $\vec{v}$ components separately, this formulation is equivalent to when $W$ is diagonal in Equation 6.8.

#### 6.2.6.2 Coupled RTide

A coupled RTide model is readily constructed by simple modification of the output layer. Here, the "Adder" 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 6.7.

#### 6.2.6.3 SIREN RTide

Fast-moving tidal currents often exhibit significant fractions of energy at high-frequencies (10). 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 (180). While these networks are more expressive, they can also be more difficult to train. The activation is defined as:

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

where $\omega_0$ 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. 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.

#### 6.2.6.4 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 dataset through backpropagation and gradient descent. The learning rate, which controls the size of each update to the model's weights during training, is initially set to $1 \times 10^{-4}$ 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 (33), 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.

### 6.2.7 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 (105). Gravitational input functions are precomputed for a full year at hourly intervals (8760 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 (63).

First, using the estimated amplitude mean $\mu$ and standard deviation $\sigma^2$ outputted by VTide from a pilot analysis of the RTide prediction, we compute the SNR of each constituents such that

$$SNR = \frac{\mu}{\sigma^2}. \tag{6.15}$$

Here, we utilize all constituents which are resolvable from a 1-year reference series using a standard Rayleigh criterion. As with other tidal analysis programs (45), 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. Even after thresholding, we encounter difficulty when using constituents within 0.5 and 2.0 SNR. While the variational Bayes estimator was shown to yield superior parameter uncertainty estimates to standard least-squares estimates, uncertainty can still be underestimated

(105; 63).

We overcome this difficulty by introducing a principled Bayesian criterion for constituent selection based on Bayes factors (181). 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

$$BF_{10} = \frac{p(D \mid M_1)}{p(D \mid M_0)}. \tag{6.16}$$

A value of $BF_{10} > 1$ indicates evidence in favour of model $M_1$, while $BF_{10} < 1$ favours 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 more crude approximation of the log of the Bayes factor due to the lack of consideration of the model prior distribution (181). 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

$$ELBO = F\left(p(\theta|Y), q(\theta|Y)\right) + \mathrm{KL}\left(p(\theta|Y), q(\theta|Y)\right). \tag{6.17}$$

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 (68). Psuedo code for the iterative algorithm is given in Algorithm 1.

Since the gravitational inputs are precomputed, the added computational cost consists of (i) the forward pass through the network to generate the predictions, and (ii) 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 6.3.2 and 6.3.3.

**Algorithm 1** Bayesian Constituent Selection via ELBO Improvement

```
constituent_list ← Rayleigh criterion constituents for full year
Run VB model on RTide prediction using constituent_list
best_ELBO ← $\text{ELBO}_{\text{Rayleigh}}$
for each constituent with SNR ≥ 0.5 in non_rayleigh_constituents do
    Run VB model on RTide prediction with constituent_list + constituent
    if $\frac{\text{ELBO}_{\text{new}}}{\text{best\_ELBO}} > 1.01$ then
        best_ELBO ← $\text{ELBO}_{\text{new}}$
        constituent_list ← constituent_list ∪ {constituent}
    end if
end for
```

### 6.2.8 Harmonic Analysis

The Harmonic Analysis described in this chapter uses a Python adaptation of the MATLAB Harmonic Analysis program UTide (34) 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 (10).

## 6.3 Validation

### 6.3.1 Simulated Comparison

In this section, we test the capabilities of the proposed architectures on simulated data from the Pentland Firth, Scotland. 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 (85; 182) to produce a depth-averaged simulation of the currents at the Pentland Firth. Simulation data is provided for the period from 2014-01-01 00:00 UTC to 2015-01-01 00:00 UTC at a sampling interval of 300 seconds. The model is forced with constituents K1, K2, M2, MU2, N2, NU2, O1, and S2 at the ocean boundary. Comparisons are made at three candidate sites characterized by increasing amounts of nonlinearity. A map of the site is provided in the following chapter in Figure 7.4

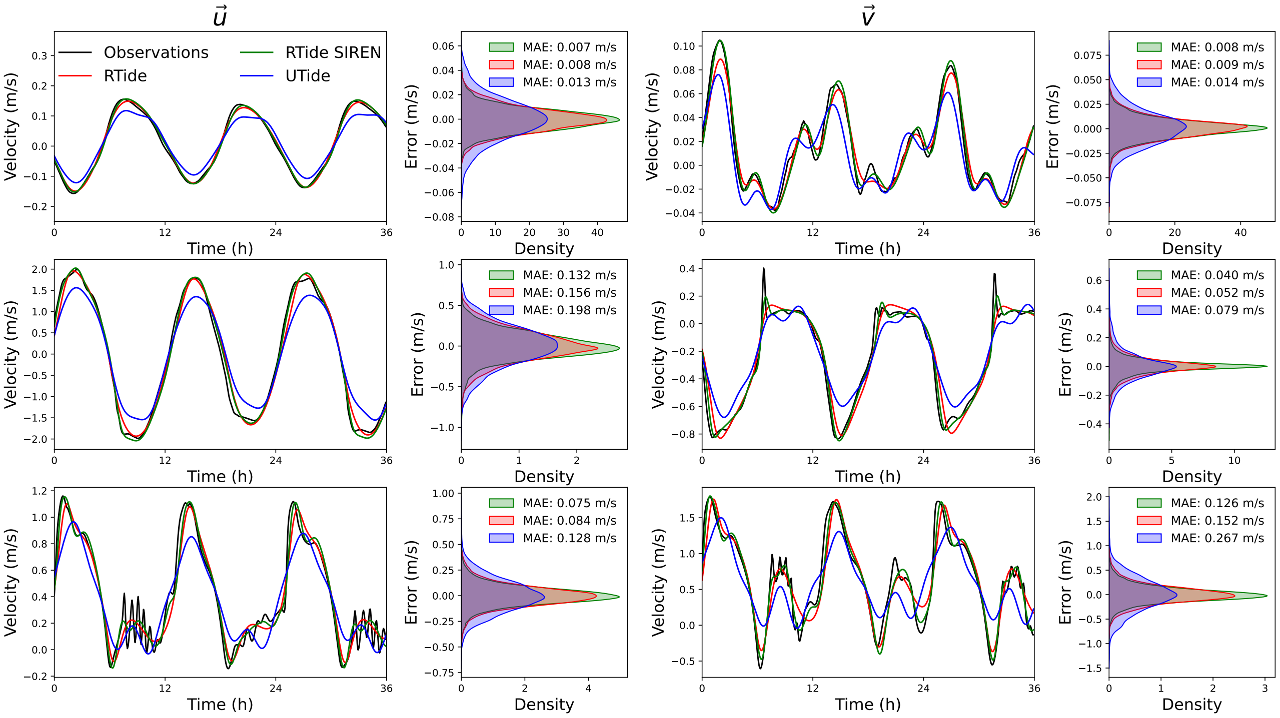

Figure 6.2: **Comparisons of UTide and RTide model's out of sample performance when trained on 180 days of simulated data.** Predictions are of the first 1.5 days of the testing dataset (2014-07-04 22:40 to 2014-07-06 10:35). Kernel density estimates of the residual error are computed from the subsequent 180 days of test set predictions. Top, middle, and bottom rows correspond to locations A, B, and C respectively with coordinates provided in Supplementary Table 1.

Visualizations of the model predictions of the simulated currents are given in Figure D.1. Inspection of the kernel density estimate (KDE) of the residuals (computed over the entire 180 day 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 6.3.2. 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.

### 6.3.2 On the limits of predictability

Most discussions of tidal current predictability have naturally focused on using harmonic analysis (183; 10). 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 6.3. 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

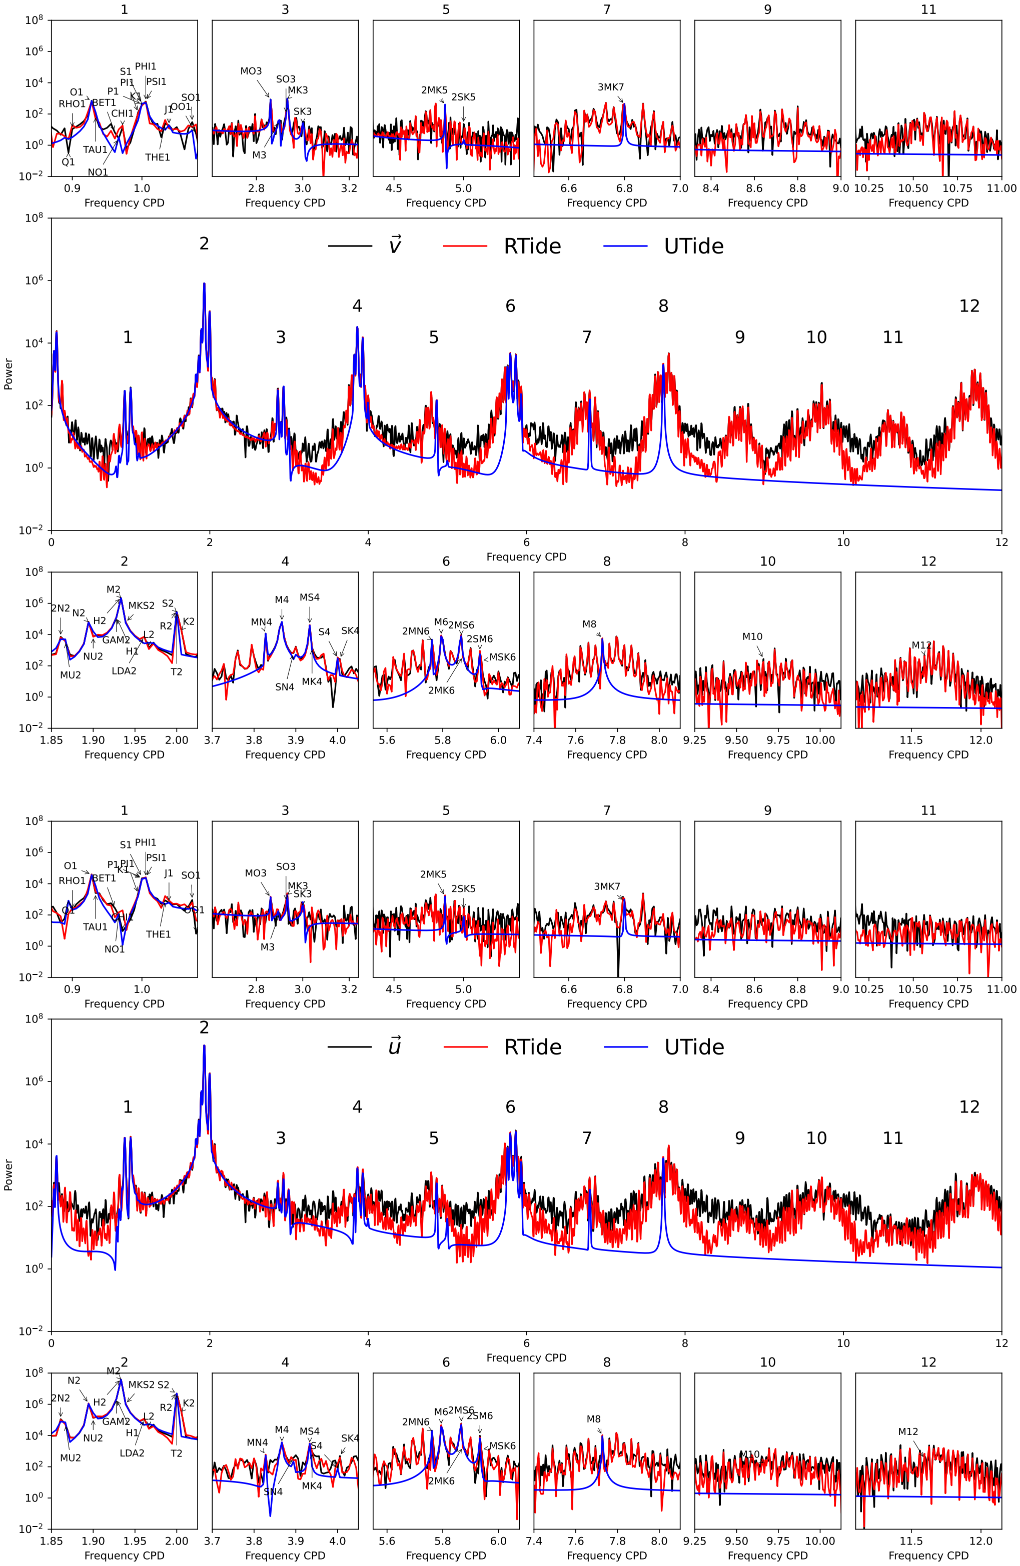

Figure 6.3: **Comparisons of predicted and observed power spectra for the simulated data between RTide and UTide.** The top panel shows the $v$-component and the bottom panel shows the $u$-component. Models are trained on 180 days of data, and PSD estimates are obtained for the 180-day test set. Smaller panels are numbered according to the species they correspond to in the central plot. All 67 standard UTide harmonic constituents are labeled.

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. 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.

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 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 6.3 is provided in the Supplementary Materials.

As has long been understood, the ability to resolve such features is also dependent on the noise of the data (1). With regard to currents, noise is often greater than for water-level measurements due to the nature of the measurement devices employed.

A comprehensive description of these sources is given in (10). We here consider how the RTide and UTide models perform in harmonic constituent estimation under varying amounts of white-Gaussian noise (WGN). Here, we utilize the harmonic equivalent response model procedure defined in Section 6.2.7. Figure 6.4 evaluates the accuracy of the RTide and UTide harmonic constituents in terms of root-sum-of-squares (RSS) error as the added WGN is increased on the Pentland Firth simulation data. The RSS in this context is defined as the sum of RMS errors across all constituents resolvable based on a Rayleigh criterion of 1. Similar behavior can be seen for both the 30-day and 180-day models. In the low-noise regime (below the ADCP standard error range), the pure HA will yield slightly more accurate results. Within the standard error range (3.9%-8.9% of the standard deviation of the flow velocity) the difference between methods is not significant. However, as the noise increases beyond this threshold the two methods diverge. Even though the WGN does not violate the underlying assumptions of least-squares, the elevated noise floor reduces the ability to resolve smaller constituents. In contrast, the response approach effectively 'filters' this noise enabling more robust 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 (44; 45) to assess whether or not that tide can be treated as stationary. Section 6.4 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 (17).

### 6.3.3 Tidal Energy Resource Assessment and Prediction

There is considerable economic interest in exploiting the energy of fast-moving tidal flows (85). Due to the significant spatial variability of tidal currents (184), it is necessary to sample multiple locations when conducting tidal energy resource as-

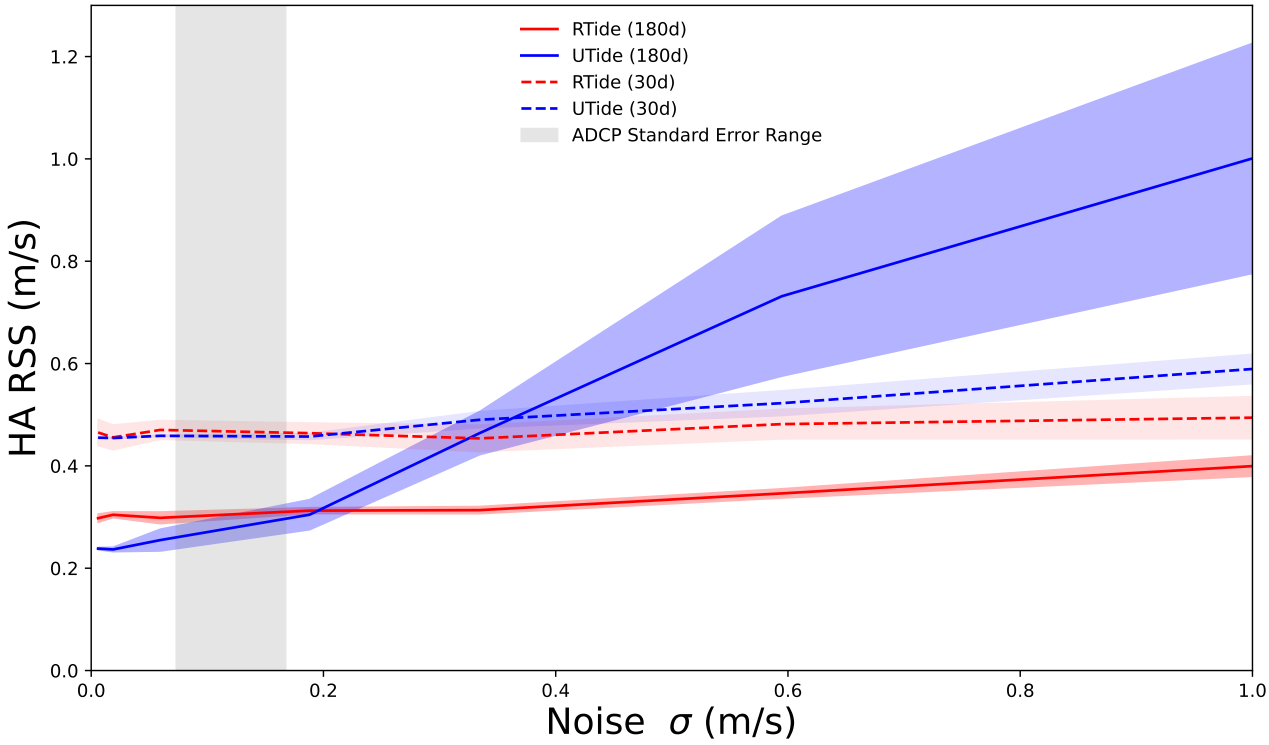

Figure 6.4: **Comparison of RTide and UTide derived harmonic constituent root-sum-of-squares (RSS) error as a function of WGN.** To enable a fair comparison, RSS is estimated only for the constituents resolvable using a Rayleigh criterion of 1. Dashed and solid lines correspond to models trained on 30-days and 180-days 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 (7).

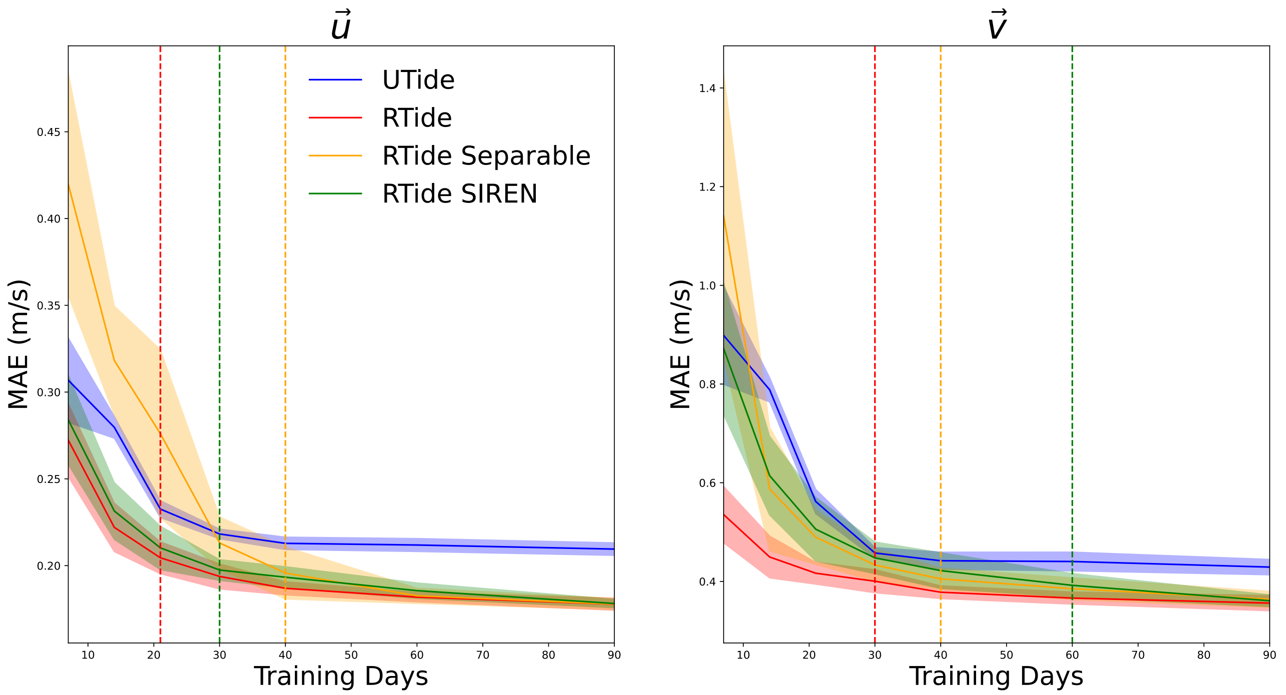

Figure 6.5: **Comparisons of UTide and RTide model's out of sample performance on 2 years of Meygen data when trained on reference series of increasing duration.** Solid lines denote mean values over 10 trials on different training data. Errors bars represent one standard deviation. Vertical dashed lines mark where the mean of an RTide model exceeds the performance of UTide trained on 180 days of data.

sessments. 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 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 (185). Presently, this is the largest tidal stream turbine deployment in the world. Data is taken from 2019-11-01 00:00 UTC to 2021-12-28 00:00 with measurements sampled at 600-second intervals without smoothing. We compare the performance of Response and HA when trained on data series ranging in length from 7 days 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 year 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 6.5.

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 the UTide, this constitutes a 13% and 41% reduction in MAE after only 7-days. 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 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 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.

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 6.3.2, 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 (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

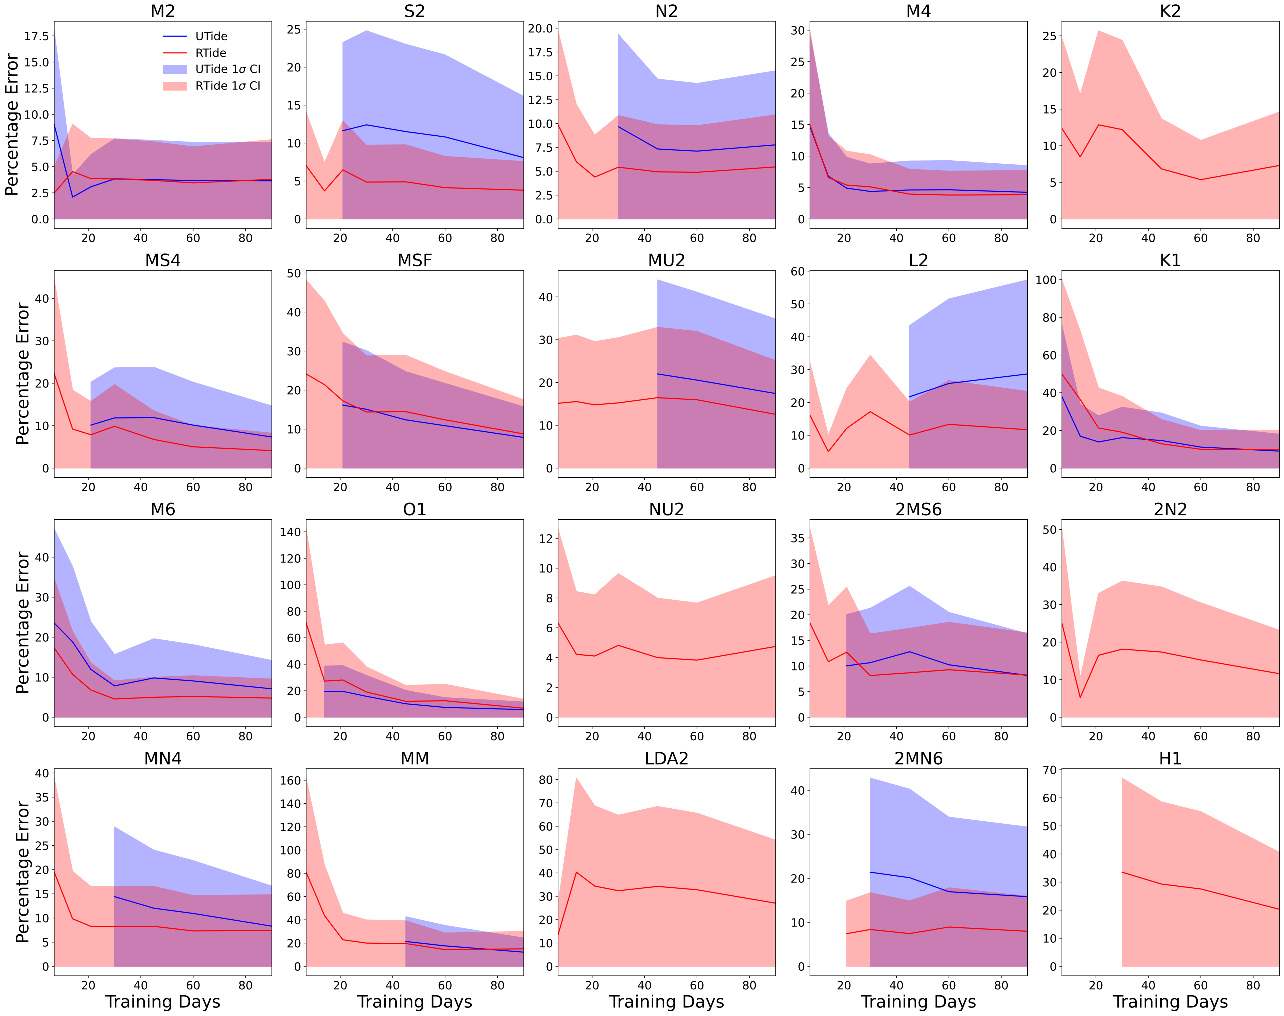

Figure 6.6: **Comparison of RTide and UTide derived harmonic constituent percentage error from Meygen data. (Equivalent to Figure 6.5)** Total percentage error is computed as the Euclidian norm of the constituent RMS error for major and minor ellipse parameters. Solid lines denote mean values over 10 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.

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 6.6, 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-years 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 6.6). Only the 20 largest constituents are shown for clarity. Analysis of individual constituents shows a similar trend to Figure 6.5 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 degrees 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-days. 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

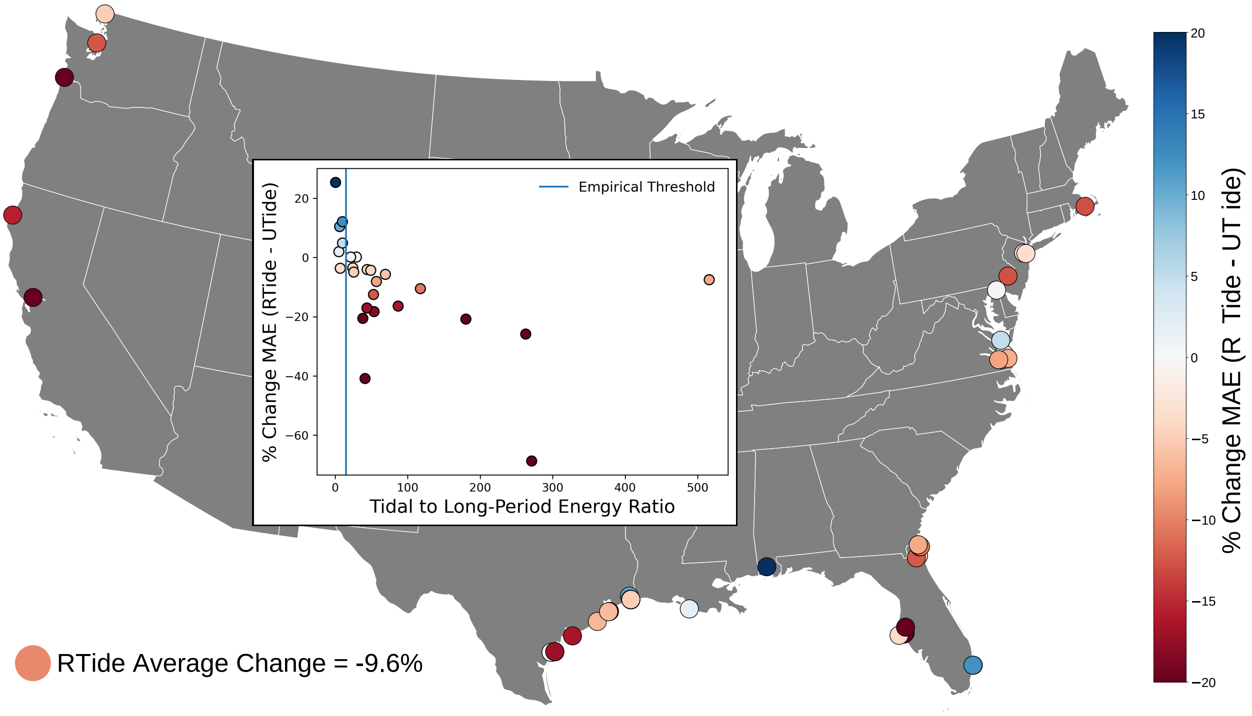

Figure 6.7: **Comparison of RTide and UTide prediction error on 40 active NOAA current stations.** Gauges are shaded according to the relative percent change in MAE between RTide and UTide. Red indicates a lower RTide MAE, and blue a higher MAE e.g. UTide is better. A single MAE value is estimated by taking the Euclidean norm of the MAE of the $\hat{u}$ and $\hat{v}$ predictions. 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.

sacrificed, and can be enhanced, when applying a response based analysis.

## 6.4 Operational Forecasting

While the proposed approach yields the greatest improvements for strongly nonlinear tidal currents, we here look to test the model's viability for operational forecasting. 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 6.2.8). 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 6.7. 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 (5; 16; 41). 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 6.3.2, 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. As shown, this variability is readily captured using our non-parametric response approach.

It is worth considering where the response approach is performing worse than the conventional HA. The inset plot within Figure 6.7 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 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 (45). 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 (15), or consider including such forcing directly into RTide as described below.

## 6.5 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 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 (41). 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 the present manuscript, the same, default, set of parameters has been utilized. In this vein, we argue that our non-parametric and data-driven approach, as implemented in RTide, carries the same "point-and-shoot" capabilities of the harmonic method, enabling the benefits of the response approach to be felt by non-experts which is key for usage in industry.

### 6.5.1 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 (15; 130). 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 (15), or include non-stationary forcing as additional inputs into RTide, as will be described below. Other approaches exist for studying non-stationary forcing such as wavelet-based approaches (45), and even spline-based methods (130). 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.

### 6.5.2 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-minute to completely run an RTide model on a single 30 day current reference series (3600 measurements) using a single Apple M3 CPU. Figure 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 (17), one can also apply *post hoc* model explainer tools such as SHAP (6). 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.

With regard to non-tidal forcing, the response approach offers distinct advantages (16). 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 (17), 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. The nature of the interactions of this forcing with the astronomical tide is thus 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 (15). 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. Separate work using this approach has also shown that similar improvements are achieved over operational numerical models (186). 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 (187). 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.

## 6.6 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. 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. All response models can be implemented using just three lines of code in the open-source Python package RTide.

## Open Research Section

All data and code needed to replicate the synthetic results are included in the RTide GitHub repository (`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.

## 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

$$\mathrm{MAE}(\vec{u}) = \frac{1}{N}\sum_{i=1}^{N}|u_i - \hat{u}_i| \tag{6.18}$$

$$\mathrm{MAE}(\vec{v}) = \frac{1}{N}\sum_{i=1}^{N}|v_i - \hat{v}_i| \tag{6.19}$$

   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

$$MAE_E = \sqrt{MAE(\hat{u})^2 + MAE(\hat{v})^2} \tag{6.20}$$

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

$$\mathrm{RMS}_k(\vec{u}) = \sqrt{\overline{[(A_u^{\mathrm{model}}\sin(\omega_k t) + B_u^{\mathrm{model}}\cos(\omega_k t)) - (A_u^{\mathrm{true}}\sin(\omega_k t) + B_u^{\mathrm{true}}\cos(\omega_k t))]^2}} \tag{6.21}$$

$$\mathrm{RMS}_k(\vec{v}) = \sqrt{\overline{[(A_v^{\mathrm{model}}\sin(\omega_k t) + B_v^{\mathrm{model}}\cos(\omega_k t)) - (A_v^{\mathrm{true}}\sin(\omega_k t) + B_v^{\mathrm{true}}\cos(\omega_k t))]^2}} \tag{6.22}$$

   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

$$\text{Total } RMS_k = \sqrt{RMS_k(\vec{u})^2 + RMS_k(\vec{v})^2} \tag{6.23}$$

# 7 A Hybrid Model for Online Short-Term Tidal Energy Forecasting

# Abstract

A hybrid model is proposed for the short-term online prediction of tidal currents. The harmonic residual analysis (HRA) model is designed to augment the numerical schemes employed by tidal energy installations by forecasting the residual error of existing methods. Using a combination of techniques from Information and Fractal Theory, a novel component selection criterion for singular spectrum analysis (SSA) is used to remove true noise from the residual time series and to decompose the signal into components that are appropriate for linear-recurrent forecasting (LRF) and high order fuzzy time series (HOFTS) respectively. The performance of the HRA method is evaluated using a combination of simulated and real data from sites in the United Kingdom and the United States. Results demonstrate the model's viability for 6-minute and 1-hour forecast horizons across sites exhibiting variable degrees of non-linearity. Empirical analysis of the resultant tidal energy forecast verifies the superior accuracy and reliability of the HRA method when compared with existing numerical schemes. Simulated data from three sites at the Pentland Firth, UK is also provided to facilitate further study of the site's power generation characteristics and to allow for direct model performance comparisons.

## 7.1 Introduction

Tidal energy is a promising form of renewable energy. Unlike alternative modes of renewable power, it is highly predictable as it is primarily driven by the relative movements of the earth, moon, and sun which can be known indefinitely far in advance. This predictability has significant advantages for grid integration (188) as well as helping to bridge periods of low production from other renewables such as wind and solar. Tidal energy extraction may be done using either an impoundment (tidal barrage) or free standing turbines (tidal stream) (189). The present paper focuses on the second of these.

While tidal currents are highly predictable in the long-term, short-term predictions (real-time to 1 hour) suffer from the stochastic and non-linear behavior introduced by turbulence, waves, bathymetric interactions, and meteorological forcing (wind stress, atmospheric pressure, etc.)(190). The resultant variation in power production and forecast accuracy imposes severe challenges to the efficient dispatching of grid resources. Accurate short-term forecasts are critical to facilitate the economical implementation of tidal energy systems and can reduce the reliance on expensive energy storage systems (191). Due to the cubic relationship between current velocity and the power produced by tidal stream turbines, small errors in current predictions can cause significant issues for tidal energy practitioners. Hence, it is imperative that high-accuracy forecasting tools are developed. Due to the significant physical and economic implications of power production, tidal energy installations primarily rely on less accurate but reliable physics-based forecasting. As outlined in Section 2.1, while many improvements have been made to tidal current forecasting using machine learning, the lack of a physical basis has inhibited the widespread adoption of these methods. Furthermore, prior models, as discussed in Section 2.1, have been developed in an offline fashion, meaning the model is pretrained on historical data and deployed. Online machine learning models continually update as data becomes available and can therefore adapt to localized changes in the underlying data distribution (192). This is critical to account for both the aperiodic behavior

| Author | Method | Location | Forecast Horizon | Type | Online | Train Time |
|---|---|---|---|---|---|---|
| (193) Fard and Su | DWT + m-SVR | Bay of Fundy, Canada | 10 minute | Data-Preprocessing + Kernel (aggregate) | No | Minutes-Hours |
| (194) Safari et al. | EEMD + LSSVM | Florida, United States | 6 minute | Data-Preprocessing + Kernel | No | Minutes-Hours |
| (195) Sarkar et al. | GP | Pentland Firth, UK (SIM) | 1 hour to 1 week | Probabilistic | No | Hours |
| (125) Sarkar et al. | GP | United States* | 12 hour | Probabilistic | No | Hours |
| (196) Safari et al. | AAQR + NPI | North America | 6,10 minute,1 hour (single-step) | Probabilistic | No | Hours |
| (197) Zhang et al. | DA-LSSVM | Straits of North America | 6 minute, 1 hour (single-step) | Optimization + Kernel | No | Minutes-Hours |
| (198) Lee and Jeng** | ANN | Taiwan* | 1 Hour,1 Year | Deep Learning | No | Hours |
| (199) Chen et al.** | DB-WT + ANN | Taiwan + South China Sea | 1-5 year(s) | Data Preprocessing + Deep Learning | No | Minutes-Hours |
| (200) Qiao et al. | GA -BPNN | Zhoushan, China | 10 minute | Optimization + Deep Learning | No | Minutes-Hours |
| (201) Aly and El-Hawary | ANN +FLSM | Bay of Fundy, Canada | 10 minute | Statistical + Deep Learning | No | Minutes-Hours |
| (202) Aly | WNN + ANN + FLSM | Bay of Fundy, Canada | 10 minute | Deep Learning | No | Minutes-Hours |
| (203) Remya et al. | EOF + GA | Numerical (SIM) | 1 hour-24 hour | Data Preprocessing +Genetic | No | Minutes-Hours |
| (204) Fard | ARIMA + SVR | Bay of Funday, Canada | 10 minute | Optimization +Linear + Kernel | No | Minutes-Hours |
| (205) Qian et al. | H-ELM + LSTM | Zhejiang, China | 5 minute | Deep Learning | No | 24 minutes*** |
| (206) Yin and Wang | OS-ELM + IGG-Fuzzy-SA | United States* | 1,3,6,12 hour (single-step) | Data Preprocessing + Statistical | Yes | .01 seconds***. |

Table 7.1: Summary of Tidal Forecasting Literature * indicates multiple ports in the same country. ** Indicates forecasts for tidal heights. *** Indicates Training/Prediction Time was provided. m- denotes ensemble/aggregate forecasting methods. The associated reference number is provided before the names of each author.

of fast-moving tidal currents and seasonal changes in flow characteristics. The aim of this paper is to develop an online forecasting method to enhance the predictive capabilities of the physical schemes currently employed by tidal energy practitioners.

## 7.2 Review of Tidal Forecasting Literature

Numerous studies have proposed methods for tidal current forecasting across short, medium, and long-term forecast horizons. We define short-term forecasts as real-time to 1-hour, medium-term forecasts as 1-hour to 1 day, and long-term forecasts as > 1 day. The two most pervasive themes in the growing tidal forecasting literature are data-preprocessing methods and imag/hybrid/combination models. Due to the quasi-deterministic nature of tidal currents, data-preprocessing methods are applied to filter noise and improve model performance. imag/hybrid/combination models have been shown to improve predictive accuracy and reliability across forecasting horizons (207; 208). These models seek to maximize the available information by leveraging the advantages of each model to produce a globally optimized forecast. While the following review exclusively looks at tidal literature, the methods used and trends presented are widely used in wind, solar, and energy load forecasting (209; 210).

After decomposing the signal into sub-harmonic components using the discrete wavelet transform (DWT), Kavousi-Fard and Abdollah forecast each sub-harmonic component using an aggregate of forecasts made by several Support Vector Machines (SVM) (193). Due to their superior performance over wavelet transforms for nonlinear and non-stationary time series, Empirical Mode Decomposition (EMD) methods have been successfully applied to tidal current prediction. Additionally, localized training was shown by (211) to improve predictive performance for nonlinear and short-term forecasts. Using an ensemble of EMDs (EEMD), (194) was able to enhance the predictive accuracy of a Least-Square Support Vector Machine through improved localized learning. The authors of (197) use the fractal characteristics of tidal current data to determine optimal training timescales. They combine this insight with a dragonfly optimization algorithm to optimize the parameter selection of the LSSVM. While SVM methods demonstrate superior generalizability over standard deep learning methods, training times are prohibitive for online applications.

Probabilistic machine learning methods have shown promise in handling the uncertainty and noise inherent to real tidal current measurements. The authors of (195) provide a method for spatiotemporal distribution modeling using a Gaussian Process (GP). Additionally, (125) demonstrate the predictive accuracy of GPs to be superior to traditional HA methods in the presence of noise and non-linearities. Furthermore, the ability to forecast probabilistic confidence intervals can help to reduce the cost of power system operation. Safari et al. develop a non-parametric interval (NPI) model with a bi-level optimization formula for extreme learning machine (BOF-ELM) to forecast confidence intervals for tidal currents and levels (196).

Due to the remarkable ability of Neural Networks to model exclusive and non-linear relationships, considerable work has gone into their application for tidal levels and currents forecasting. In 2002, Lee and Jeng demonstrated that Artificial Neural Networks (ANN) could produce higher accuracy forecasts than HA (198). These deep learning methods suffer from a number of drawbacks; training is computationally intensive, data-hungry, and highly prone to overfitting or converging to local optima. A variety of techniques have therefore been applied to combat these shortcomings.

Chen et al. employed the Daubechies Wavelet transform (DB-WT) to remove the high-frequency 'noise' components before using Artificial Neural Networks (ANN) to predict tidal heights (199). In order to prevent local optima convergence, Remyra et al. use a Genetic Algorithm (GA) to determine optimal Back Propagation Neural Network (BPNN) parameters for short-term tidal current prediction (203). In (201), tidal currents are first forecasted using a Least Squares Fourier Series Model (FLSM). The associated prediction error referred to as 'innovations' is then used as input to the ANN. The final forecast is the sum of the FLSM prediction and the ANN forecast. Long-Short-Term Memory (LSTM) models have also been used for tidal current forecasting due to their ability to capture both long and short-term dependencies. The authors of (205) develop a Hybrid Hierarchical Extreme Learning Machine (H-ELM) and LSTM method to predict the random and periodic components of multi-layer tidal currents.

Online tidal level predictions were realized in (206) using a sample pool selection scheme to optimize an online sequential extreme learning machine (OS-ELM) in conjunction with the improved Gath-Geva Fuzzy Segmentation algorithm (IGG-Fuzzy-SA) and harmonic analysis (HA). To the best of our knowledge, this is the only study that has realized online tidal prediction and only considers tidal levels. Hence, there is a clear need for an online tidal current prediction method.

Outside of the tidal literature, multiple works have looked to combine the interpretability and reliability of physical methods with machine learning for short-term forecasting tasks. The authors of (212) combined a theoretical model of solar radiation with an ANN for the task of photovoltaic power plant production forecasting. The PH-ANN model was able to learn the non-linear cloud cover dynamics missed by the solar radiation model and outperformed a stand-alone ANN. In (213), the predictive capabilities of a modified gravity model were improved using a CNN encoder for the task of metro transit passenger flow forecasting. Daily precipitation forecasts were bolstered in (214) by using XGBoost for bias correction of a numerical weather prediction model. Similarly, the authors of (215) showed that daily evapotranspiration estimation could be improved by using a particle swarm optimization algorithm (PSO) to combine and optimize two surface temperature-vegetation index (Ts-VI) triangle models.

In order to combat the limitations faced by traditional statistical and deep learning models we propose a Harmonic Residual Analysis scheme (HRA) for online short-term tidal current forecasting. The HRA is designed to be easily modified to fit into the forecasting methods currently used by practitioners. With cumulative training and forecasting times of less than 10 seconds, our algorithm is capable of real-time online forecasting. Furthermore, the SSA-LRF and HOFTS only require 3 and 5 days of consecutive training data respectively. This form of localized training allows for rapid adaptability to variable conditions. The automated component selection scheme for the SSA-LRF is able to remove noise and identify components that are best suited for the LRF and Fuzzy forecasts respectively. By incorporating the stability of physical methods with the local adaptability of the combined SSA-LRF and HOFTS methods, HRA provides reliable and high-accuracy short-term online forecasts.

## 7.3 Overview of Methods

### 7.3.1 Harmonic Analysis

Tidal currents are primarily driven by the gravitational potentials induced by movements of the Earth, Sun, and Moon and can be described by Potential Field Theory. As such, a tidal prediction model can be derived from the characteristic harmonic frequencies of the elliptical orbits of these astronomical bodies. The resultant model is defined by tidal constituents which correspond to a particular harmonic frequency. This fundamental idea is the basis of the Harmonic Analysis method whose basic procedure is as follows:
For a time series of current velocity observations $y(t)$ with time steps $t$ at fixed

intervals $\Delta t$ we can predict the tidal heights using $N$ constituents such that

$$y(t) = \sum_{k=1}^{N} (a_k^+ e^{i\omega_k (t-t_0)+i\beta_k} + (a_k^- e^{-i\omega_k (t-t_0)-i\beta_k}) + x_0 + x_1(t - t_0), \tag{7.1}$$

where $a_k^+$ and $a_k^-$ correspond to the unknown complex amplitudes for the $k^{th}$ components, $w_k$ is the angular frequency and $\beta_k$ is the corresponding astronomical argument for the $k^{th}$ components. The inclusion of the additional terms $x_0$ and $x_1$ is necessary to model the offset and trend respectively. Given that we are interested in modeling currents, $y(t)$ will have the form $y(t) = U(t) + iV(t)$ with the real and imaginary parts corresponding to the East-West and North-South Velocity components respectively. Fitting of this equation can be carried out using an Iteratively Reweighted Least-Squares Fitting procedure (IRLS). Let $\phi_m^k = w_k(t - t_0) + i\beta_k$. We can now rewrite Equation 1 in matrix form such that

$$\begin{bmatrix} e^{i\phi_1^1} & \dots & e^{i\phi_1^N} & e^{-i\phi_1^1} & \dots & e^{i\phi_N^1} & 1 & (t_1 - t_0) \\ e^{i\phi_2^1} & \dots & e^{i\phi_2^N} & e^{-i\phi_2^1} & \dots & e^{-i\phi_2^N} & 1 & (t_2 - t_0) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ e^{i\phi_M^1} & \dots & e^{i\phi_M^N} & e^{-i\phi_M^1} & \dots & e^{-i\phi_M^N} & 1 & (t_N - t_0) \end{bmatrix} \begin{bmatrix} a_1^+ \\ \vdots \\ a_N^+ \\ a_1^- \\ \vdots \\ a_N^- \\ c_0 \\ c_1 \end{bmatrix} = \begin{bmatrix} u_1 + iv_1 \\ u_2 + iv_2 \\ \vdots \\ u_N + iv_N \end{bmatrix}. \tag{7.2}$$

In order to reduce the sensitivity to non-tidal variations the IRLS method makes use of an iteratively determined weighting matrix $W$. A detailed description of the weighting matrix $W$ and how to calculate it is given in (216). The final solution is obtained using

$$a = (S^* W S)^{-1} S^* W y, \tag{7.3}$$

where $S$ is a matrix of sinusoidal basis functions evaluated at $t_i$. Given a sufficiently long time series (18.6 years) it is possible to resolve more than 500 tidal constituents using Harmonic Analysis. However, within the context of tidal current analysis, continuous time series longer than 60 days are rare. As such, it is important that we impose an automated criterion for tidal constituent selection. The Rayleigh criterion defines the minimum resolvable frequency between two constituents to be $T^{-1}$ for a time series of length $T$ (216). The Harmonic Analysis described in this paper is carried out using a Python adaptation of the MATLAB Harmonic Analysis program UTide (34) with a Rayleigh threshold $R = 1$.

### 7.3.2 Singular Spectrum Analysis (SSA)

SSA is a powerful non-parametric tool for time-series analysis and filtering (217). Similar to Principal Component Analysis (PCA) but for one-dimensional series, SSA decomposes a signal into independent oscillatory components. These components can then be grouped together and reconstructed to produce a filtered version of the input series. Common applications of SSA are noise removal, trend/periodicity extraction, and of particular interest to us: forecasting. We outline the 4 primary steps of SSA below:

1. **Embedding** The embedding procedure transforms the input time series into a sequence of lagged vectors using a sliding window. Once again consider an arbitrary time series $y(t)$, for convenience we let $y_i = y(t_i)$. Our time series can be written as $\{y_0, y_1, \dots, y_T\}$ where $T$ is the length of the time series. The embedding procedure is defined in terms of the window length $k$, where $k < \frac{T}{2}$. The number of lagged vectors $X_j$ produced is therefore given by $L = T - K + 1$ where $X_j = (y_j, \dots, y_{j+L-1})^T$ for $j = 1, 2, \dots, K$. The lagged vectors are then transformed into a 2-d Hankel Matrix of the form

$$X = \begin{bmatrix} x_1 & \dots & x_K \\ \vdots & \ddots & \vdots \\ x_L & \dots & x_T \end{bmatrix} \tag{7.4}$$

2. **Decomposition** The decomposition of our Hankel Matrix is carried out using Singular Value Decomposition (SVD). The Eigenvectors and Eigenvalues are calculated for $S = XX^T$. Since $X$ is a Hankel matrix, we recognize $XX^T$ is

symmetric and diagonal. The eigenvectors $U_i$ and corresponding eigenvectors $\lambda_i$ for $i = 1, 2, \ldots, d$ where $d = rank(X)$ are then arranged in descending order. Note $L = d$ if all eigenvalues are non-zero. For each eigenvector $U_i$, we compute the corresponding vector $V_i = X^T \times \frac{U_i}{\sqrt{\lambda_i}}$. The final SVD is the sum of all $V_i$ with principal components given by elementary matrices $E_i$ of the form

$$E_i = \sqrt{\lambda_i} U_i \times V_i^T. \tag{7.5}$$

3. **Grouping** The Grouping stage is critical to the success of the SSA and subsequent Linear Recurrent Forecast and consists of grouping the eigentriples $(\lambda_i, U_i, V_i)$ into disjoint subsets. The eigenvalues $\lambda_i$ are sorted in descending order with their magnitudes defined by their principal components' contribution to the overall signal variance $S^2$. Typical eigentriple grouping procedures partition the components into disjoint sets based on their relative contributions. Some commonly used methods are grouping the first $n$ components that account for 95% of the variance, correlation-matrix inspection, mean-based selection, and clustering (218). We demonstrate the need for an alternative method in Section 3.4 and present our unsupervised algorithm for component selection in Section 3.5.

4. **Reconstruction** Once we have grouped the elementary matrices from the SVD we recover the corresponding time series using Diagonal Averaging. Let $L^* = min(L, K)$, and $K^* = max(L, K)$. The reconstructed series $y'(t)$ is given by

$$y'(t) = \begin{cases} \frac{1}{k} & \sum_{m=1}^{k} y_{m,k-m+1} & 0 \le k \le L^* \\ \frac{1}{L^*} & \sum_{m=1}^{L^*} y_{m,k-m+1} & L^* \le k \le K^* \\ \frac{1}{T-k} & \sum_{m=k-K^*+1}^{T-k+1} y_{m,k-m+1} & K^* \le k \le T \end{cases} \tag{7.6}$$

#### 7.3.2.1 Linear Recurrent Forecasting (LRF)

Linear Recurrent Forecasting relies on the assumption that our newly reconstructed series satisfies a Linear Recurrent Relation (LRR). Through careful selection of SSA components, we are able to extract a signal that satisfies this condition. That is, our time series $y'(t) = \{y'_0, y'_1, \ldots, y'_T\}$ can be expressed as a LRR of order $d$ such that

$$y'_{i+d} = \sum_{k=1}^{d} c_k y'_{i+d-k} \quad \text{for} \quad 1 \leq i \leq T-d \tag{7.7}$$

where $c_k$ are the LRR coefficients which can be estimated using the method outlined in (219), and order $d < T$. Using this relation the Linear Recurrent Forecasting prediction $\tilde{y}$ is defined by

$$\tilde{y}(t) = \begin{cases} y'_i & i = 1, 2, \ldots, T \\ \sum_{j=1}^{N-1} c_j \tilde{y}_{i-j} & i = T+1, \ldots, T+H \end{cases} \tag{7.8}$$

where $H$ is the number of steps ahead to forecast. It can be clearly seen that the Linear Recurrent Forecast is simply a weighted sum of the previous values of the reconstructed signal. Improper SSA component selection will therefore lead to inaccurate forecasts due to our assumption that $y'$ satisfies a LRR. As such, we develop a robust method for SSA component selection in the following two sections.

### 7.3.3 Permutation Entropy (PE)

PE measures the complexity of a time series through the comparison of adjacent values. The partitioning of our time series into a set of 'ordinal patterns' and subsequent comparisons allow for the discrimination between chaotic, deterministic, and stochastic processes (220). PE has been shown to be robust to noise and invariant to other non-linear transformations. Furthermore, PE is applicable to arbitrary time series and is efficient to calculate making it an ideal candidate for the real-time selection of SSA components.

The complete time series can be partitioned into the set of ordinal patterns $\Pi = \{\pi_1, \pi_2, \ldots \pi_{n!}\}$ defined in terms of the order $n$ and the time lag $\tau$. The order $n$ determines the number of values contained within a given permutation. It follows that the set of all permutations $\Pi$ will have cardinality $n!$. The time lag $\tau$ determines the number of time steps spanned by each segment of a permutation. Due to its temporal significance, $\tau$ plays a critical role in determining the frequency characteristics of the PE (221). The relationship between the order $n$ and the time lag $\tau$ is illustrated in Appendix 7.10.

The procedure for calculating the PE can be defined in three steps:

1. Partition the input series into a sequence of sets whose length is equal to $n$.
2. Sort each set according to the relative magnitude of its components and identify its corresponding $\pi_i \in \Pi$.
3. Sum the number of sets corresponding to each $\pi_i$ and compute the probability $p(\pi_i)$ for each $\pi_i$ such that

$$p(\pi_i) = \frac{\text{Number of Occurrences of Corresponding Sets}}{\text{Total Number of Sets}}. \tag{7.9}$$

The final permutation entropy $H(n)$ is given by

$$H(n) = \sum_{i=1}^{n!} p(\pi_i) \log(p(\pi_i)). \tag{7.10}$$

It is clear the range of values $H(n)$ can take on is $0 \leq H(n) \leq \log(n)$. The normalized PE can therefore be written as $N(PE) = H(n)/\log(n)$. For the remainder of this paper, when referring to PE we will be referring to the normalized PE (i.e. $H(n)/\log(n)$).

Given $H \in [0,1]$ it remains to be seen what the physical significance of these values are within the context of our time series. There are three cases:

1. $H(n) \approx 0$ : The time series is deterministic. This corresponds to a singular $\pi_i$ being more dominant.
2. $0 < H(n) < 1$ : The time series is characterized by both deterministic and stochastic processes. This corresponds to a varied distribution of $\pi_i$s.

3. $H(n) \approx 1$ : The time series is stochastic. The distribution of $\pi_i$'s is uniform.

We elect to calculate PE using the Python package Ordpy described in detail in (222).

### 7.3.4 Complexity Analysis of SSA Components

To establish the motivation for our automated component selection procedure we look at the SSA Decomposition of the UTide Residual of data collected from Estes Head (EPT0003) located in Eastport, Maine. The Eastport data was collected from 01 July, 2000, to 30 August, 2000, with a sampling rate of 6 minutes. The PE of the components of the SSA decomposition and their relationship to their respective 'significance' levels are shown in Figure 7.1. The decomposition was carried out for a 7-day subset (1680 Samples) of the original series using a window length of $k = 540$.

It can be clearly seen that the magnitude of the SSA component's contribution to the overall variance is not necessarily an indication of its association with noise. As outlined in Section 3.2 and shown in Figure 7.1, typical SSA grouping procedures fail to account for this behavior and include true noise in our reconstruction whilst simultaneously discarding deterministic information. Hence, a principled method is needed in order to ensure the proper selection of SSA components.

### 7.3.5 Unsupervised Binary Grouping of SSA Components

Due to site-dependent temporal and frequency characteristics as well as the added degrees of freedom introduced by PE parameters $n$ and $\tau$ it is impossible to impose a general PE threshold with which to describe a signal as being either noisier or more deterministic. Thus, our approach is to empirically determine this threshold using the decomposed components' PE characteristics. We adopt a method similar to (223), using a simple Unsupervised Binary Grouping algorithm (UBG). Given two sets, UBG seeks to maximize the distance between the means of each set. The eigenvalues of the SSA decomposition do not contain any information relating to the structure or complexity of its signal. As such, we must first reconstruct each

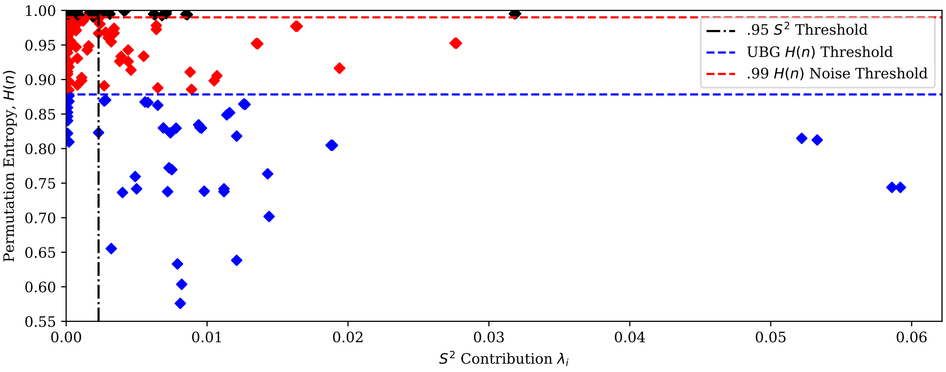

Figure 7.1: Permutation Entropy of Reconstructed SSA Components for 7-day Eastport UTide Residuals. $S^2$ Contribution refers to the eigenvector $\lambda_i$'s contribution to the overall signal variance. PE is calculated for each reconstructed signal for the length of the original series (1680 Samples). Points are coloured according to their membership to either the LRF signal, the HOFTS signal, or true noise signal denoted by blue, red, and black points respectively. The standard .95 $S^2$ threshold is shown in black. The Unsupervised Binary Grouping $H(n)$ threshold is shown in blue. The true noise threshold for components with $H(n) > .99$ is shown in red.

component to determine its PE. After computing the set $\{H_1, H_2, \dots, H_{k-1}\}$ of PEs for the reconstructed signals, we calculate the threshold $H_{T1}$ using the following Binary Grouping Procedure:

1. Let $H_{T0} = mean(\{H_1, H_2, \dots, H_{k-1}\})$
2. Form two sets, $A, B$ with $A$ being the set of components with $H(n) < H_{T0}$ and $B$ the set of components with $H(n) \geq H_{T0}$ and compute the mean values $\mu_A, \mu_B$ for each.
3. $H_{T1} = (\mu_A + \mu_B)/2$.
4. Continue until $|H_{T0} - H_{T1}| < \epsilon$, where $\epsilon$ is the termination condition.

During our testing, $\epsilon = .001$ proved to be more than sufficient for our application. While noise adaptive thresholding is effective at separating the low and high-frequency components from the reconstructed signal it does not guarantee the signal exhibits sufficient self-similarity for LRF. As such, we adopt the Hurst Exponent as a measure of the series' self-similarity. A complete explanation of the Hurst exponent and how it is calculated is given in (224). If a series has a Hurst exponent greater than 0.5 then it contains self-similarity and can therefore be used for LRF. If a series has a Hurst index of 0.5 or less then it is non-independent or randomly independent and cannot be used for LRF. In this case, the SSA is simply used to remove noise from the signal and the combined low and high-frequency signals are provided to the Fuzzy Time Series method outlined in the next section.

### 7.3.6 Fuzzy Time-Series Forecasting

Interest in Fuzzy Time Series (FTS) has grown dramatically in recent years due to the advantages FTS provides over conventional methods. Namely, FTS requires significantly less data than statistical methods (225). FTS makes use of linguistic variables to form simple logical relationships and generate subsequent forecasts. As such, FTS forecasting is efficient, interpretable, and adaptable. FTS is built on fuzzy set logic. Traditional set logic states that elements must belong to single sets, membership $x$ is therefore defined by $x \in \{0,1\}$. Fuzzy sets allow for simultaneous set membership with membership being defined for the range $x \in [0,1]$. The basic procedure for FTS Forecasting can be defined as follows:

1. **Universe of Discourse** We define the Universe of Discourse, $\mathcal{D}$, for our time series $y(t)$ as $\mathcal{D} = [\min(y(t)), \max(y(t))]$ where $\mathcal{D} = \{u_1, u_2, \ldots, u_n\}$.
2. **Partition $\mathcal{D}$** We can partition the Universe of Discourse into a set of overlapping intervals defined by the fuzzy sets $f_j$. Let $f_A$ be an arbitrary fuzzy set membership function such that $f_A \to [0,1]$. The fuzzy set $A$ is given by

   $$A = \frac{f_A(u_1)}{u_1} + \frac{f_A(u_2)}{u_2} + \cdots + \frac{f_A(u_n)}{u_n}. \tag{7.11}$$

   The resultant FTS $F(t)$ is defined by the set of all fuzzy sets $f_j$ for $j = A, B, \ldots,$ defined on the Universe of Discourse $\mathcal{D}$. There are two primary types of membership functions: even spacing, and uneven spacing. [2] provides a more detailed description of partitioning methods and their applications.
3. **Create Fuzzy Logical Relationships** We assume that there exists a relationship $R(t-1,t)$ between consecutive values $F(t-1)$ and $F(t)$. This implies that $F(t) = F(t-1) \circ R(t-1,1)$ where $\circ$ is an arbitrary composition operator. Thus the Fuzzy Logical Relationship (FLR) is given by $F(t-1) \to f(t)$. We

can extend this treatment to include multiple FLRs by assuming $F(t)$ is caused by the $N$ preceding terms. It follows that $F(t-N),\ldots,F(t-1) \rightarrow F(t)$.

4. **Forecast** Using the FLRs defined in the previous section we can compute future values by identifying their corresponding FLR. Forecasted fuzzy values are de-fuzzified based on the fuzzy set membership functions defined in Step 2.

While FTS can be applied to arbitrary time series, we found it necessary to difference our time series prior to applying FTS in order to remove non-stationarity and improve forecasting accuracy. A number of methods have been proposed for FTS; we make use of Chen's method which leverages High Order Fuzzy Time Series (HOFTS). A full description of Chen's method is given in (226). FTS Forecasts in this paper are implemented using the pyFTS toolbox (227).

## 7.4 Harmonic Residual Analysis Forecasting Scheme

Based on the methods above, we present an Online Harmonic Residual Analysis Forecasting scheme that utilizes HA-SSA-UBG-LRF-FTS which we refer to as (HRA).

The basic procedure for Real-Time Online prediction can be summarized as follows:

1. Perform HA on the input time series, generate a forecast, and compute the residuals between the HA and the observed data.
2. HA residuals for the previous 3 days of observations are fed into the SSA. The series is embedded into the Hankel Trajectory Matrix and decomposed using SVD.
3. Individual signals are reconstructed for each eigentriple. The PE of the resultant signals is computed and stored. UBG is applied to determine the optimal PE threshold with which to define our signal and noise sets.
4. Components with $PE > .99$ are removed. The remaining components are partitioned into two sets corresponding to low and high PE.

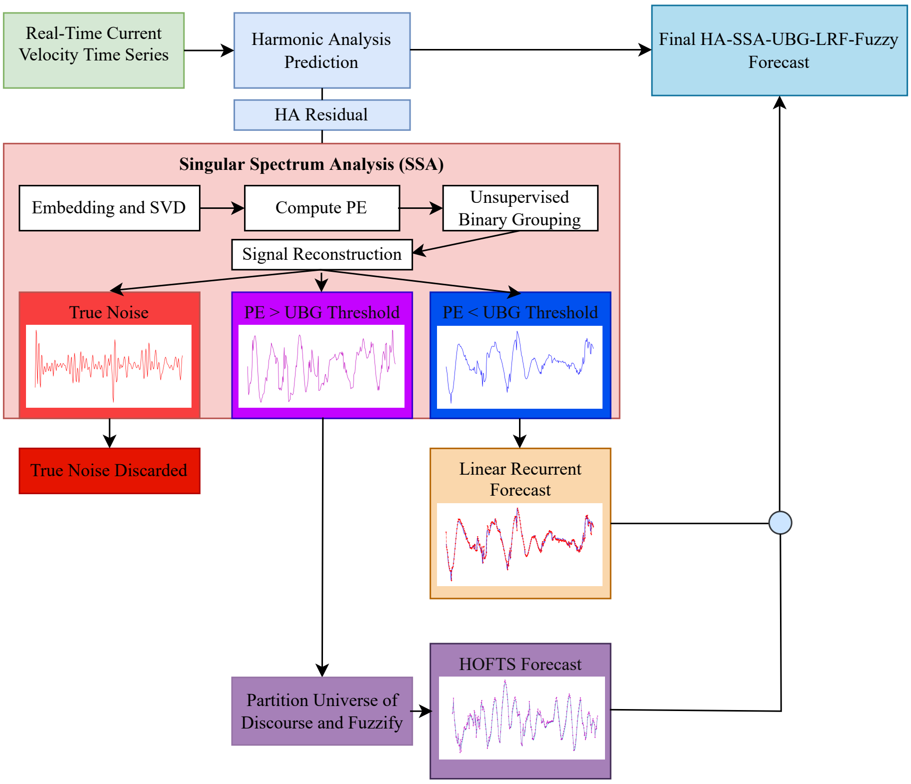

Figure 7.2: HRA Model Architecture. HA = Harmonic Analysis, SVD = Singular Value Decomposition, PE = Permutation Entropy, UBG = Unsupervised Binary Grouping, LRF = LRF, HOFTS = High-Order Fuzzy Time Series.

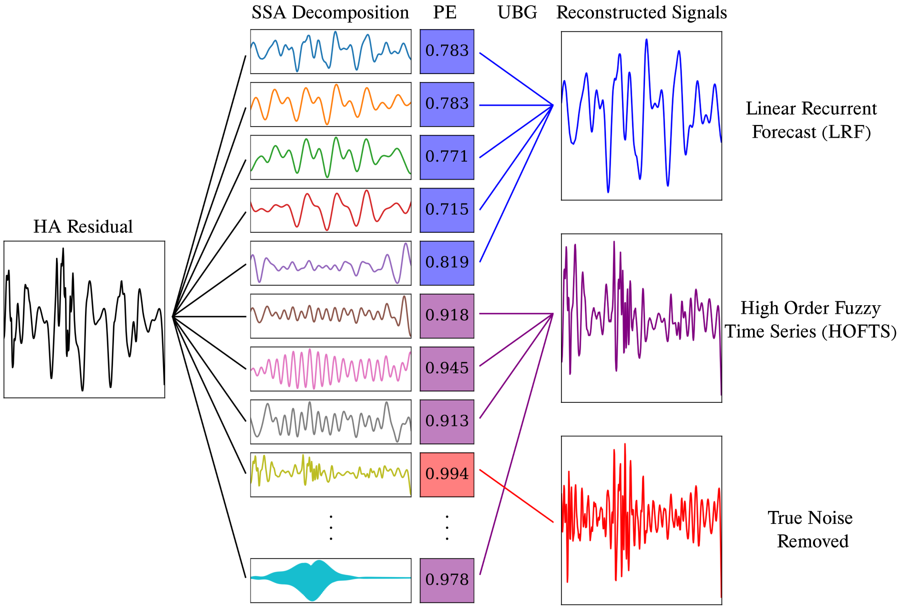

Figure 7.3: SSA decomposition and automated grouping explained. Components are shown in descending order of their contribution to the signal $S^2$ and scaled for visualization.

5. The high PE set is reconstructed and set aside to be predicted using HOFTS.
6. The components of the low PE set are reconstructed and the moving Hurst exponent is calculated. If the Hurst exponent is greater than 0.6 then the signal demonstrates sufficient autocorrelation and is forecasted using LRF. If the Hurst exponent is less than 0.6 then the signal is inappropriate for use in LRF and is combined with the high PE signal in the next step.
7. The high PE signal is then presented to the HOFTS. The data is differenced, partitioned using an entropy partitioner, and then predicted using the Chen Method whose procedure is outlined in Section 3.6. Due to the low computational cost of HOFTS, we make use of the entire set of historical residuals. As such, the performance of the HOFTS will improve as the duration of the Online implementation increases.
8. The final HRA Forecast is the ensemble of predictions from the Harmonic Analysis, Linear Recurrent Forecast, and Fuzzy Time Series.

Due to the residual-based construction of our prediction scheme, two weeks of continuous observations are necessary in order to realize real-time Online predictions. This is the first time SSA-LRF and HOFTS have been used to forecast tidal currents. The proposed model is similar in structure to the model outlined in (228), however, we have made a number of critical changes to their proposed architecture. First, the automated component selection for SSA ensures that the reconstructed signal is suitable for LRF. Using the novel selection criterion based on information and fractal theory, our model is able to recognize and adapt to local anomalies. In this way, the HRA method is insulated from the erroneous and unpredictable errors faced by other ML methods. Second, we find it to be inappropriate to use the residuals of the LRF in our HOFTS prediction. Due to the fact that component grouping is not necessarily uniform between decompositions, LRF residuals may not be consistent between decompositions and can lead to inaccurate forecasts. Additionally, by using ML methods to bolster the performance of physical schemes, the resultant HRA model retains the reliability demanded by tidal energy practitioners and combines it with the high accuracy of ML.

### 7.4.1 A Note On Data Leakage

Data leakage is a critical consideration when decomposing time series and can result in erroneous predictions as well as model testing not being rigorous (229). A brief overview of the two main types of leakages and how HRA avoids them is given below.

1. ***Feature Leakage*** *results from the inclusion of data that will not be available to the model in testing.* HRA predictions are computed using only the past values of the univariate reference series and are therefore uncontaminated by feature leakage.
2. ***Train-Test Leakage*** *occurs when there is an overlap between the train and test sets for model training or decomposition. This can occur as a result of random sampling, or if decomposition is carried out on the entire time series prior to separation (one-time decomposition)*(229). HRA avoids train-test leakage through temporal partitioning (partitioning before and after a given timestamp) prior to the HRA algorithm being run. In this way, HA and the SSA decomposition are only run on the train set and receive no information or overlap with the test series.

## 7.5 Results

### 7.5.1 Pentland Simulation

The Pentland Firth, which connects the Atlantic Ocean to the North Sea, is characterized by exceptionally fast tidal currents, thus making it a high-interest site for tidal energy development. We have adopted the numerical scheme proposed and validated in (85) to simulate the currents at the Pentland Firth. We first benchmark our model using this simulated data in order to test our model's ability to capture severe non-linear current interactions without the presence of measurement noise or meteorological effects. The simulation models the tidal flows in the Pentland Firth by solving the shallow water equations using the Discontinuous Galerkin (DG) method (230). The depth-averaged model is forced with the tidal constituents

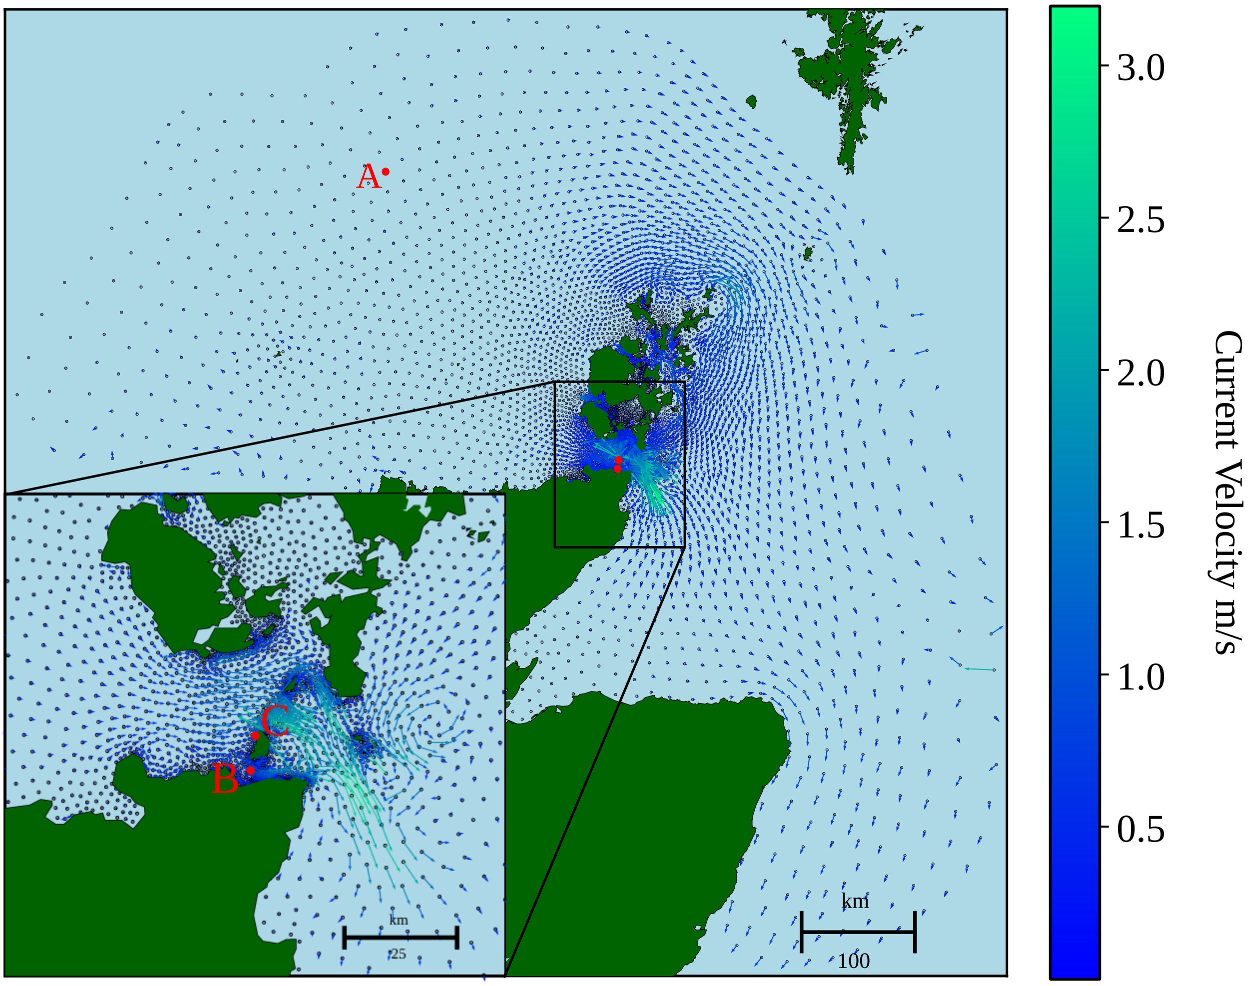

Figure 7.4: Map of the Pentland Firth with node locations and velocity field overlaid. The depth-averaged simulation contains 5200 nodes with grid resolution varying from 150m in the Pentland Firth to 20km beyond the continental shelf.

K1, K2, M2, MU2, N2, NU2, O1, and S2 and shows reasonable agreement with field data at various locations.

A plot of the Pentland Firth, node locations, and the simulated currents is shown in Figure 7.4. The time series consists of simulated observations from January 1st, 2014 through October 31st, 2014 at a sampling rate of 5 minutes. We present a case study of three nodes characterized by increasing levels of non-linearity. Table D.1 provides the coordinates for each node and numerical designations.

#### 7.5.1.1 Experimental Setup

To test our model's forecasting performance we benchmark using a combination of simulated and real data. We provide case studies of sites characterized by increasing degrees of non-linearity. Model performance is tested for different sampling intervals and the power predictive capabilities are evaluated. We assess multiple aspects of our model's predictive capabilities using Root-Mean Squared Error (RMSE), Mean Squared Error (MSE), Mean Absolute Error (MAE), Maximum Absolute Error (Max Error), and R-Squared ($R^2$). The definition of each statistical measure is given in 7.7.1. Additionally, the SSA parameters used for each site can be found in 7.7.5, Table 7.8.

Due to the fact that tidal energy site stakeholders demand a physical basis for forecasting, the focus of this chapter is not on pure Machine learning methods. A comprehensive set of comparisons between the HRA method and current state-of-the-art ML methods (RF, SVR, ANN, and LSTM) can be found in 7.7.4 for both real (Table 7.7) and simulated data (Table 7.6). It should be noted that HRA outperformed or tied the best-performing ML methods for 26/30 and 19/20 performance metrics for the simulated and real data respectively. Furthermore, it can also be seen that the pure ML methods have significant variability in performance accuracy across sites. A description of the model architectures and hyperparameters for each method can also be found in 7.7.4.1.

#### 7.5.1.2 HRA Model Performance (Simulated Data)

HA predictions were computed using 15 days of reference data. SSA-LRFs were then run for 7 days using the computed HA residuals and a training-set size of 72 hours (720 observations). The Hurst exponents for the 72-hour training sets are greater than 0.7 indicating the decomposed signal has a high level of long-range correlation and is thus suitable for LRF. The high-PE signal was then provided to

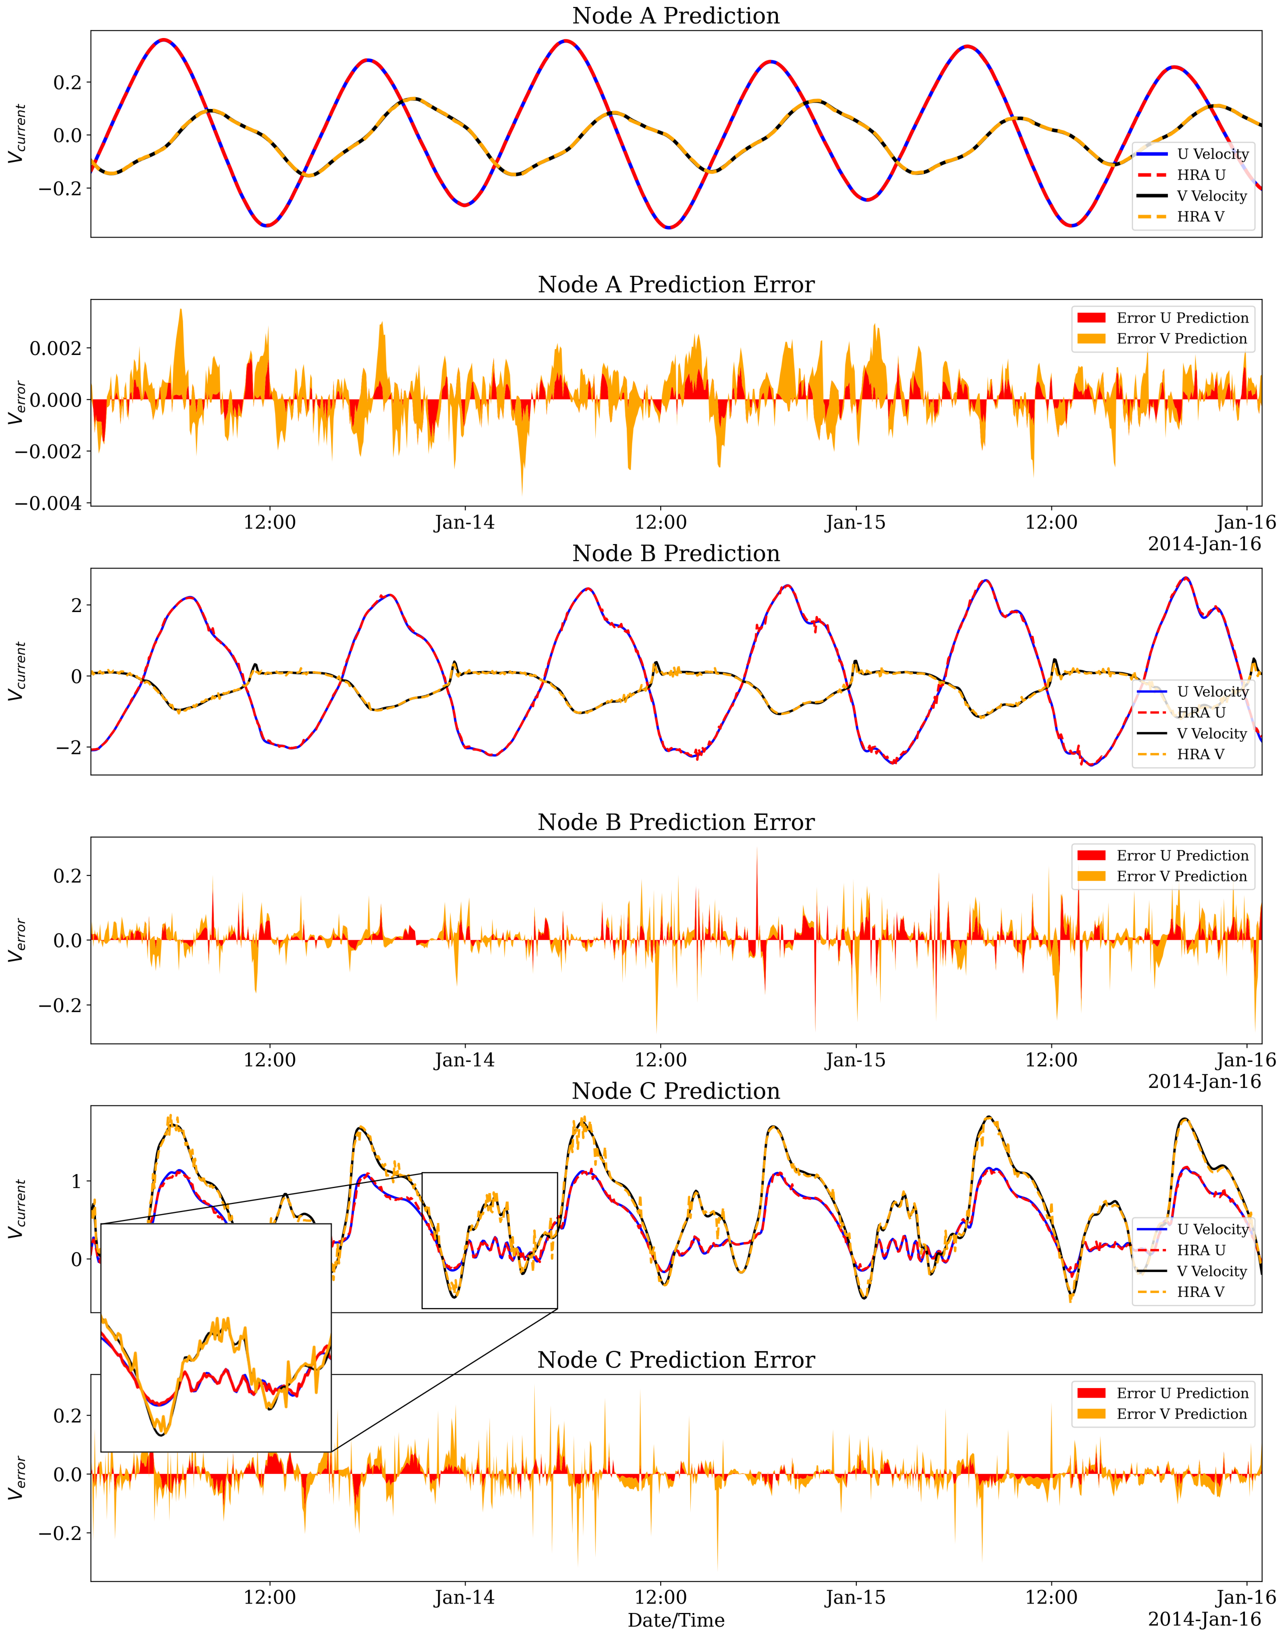

Figure 7.5: 72 hours of 6 Minute Online Forecasts and associated prediction error for Pentland Simulation Data. Note the current velocity error axes are scaled for each node.

| Node | Prediction Data | RMSE | MSE | MAE | Max Error | $r^2$ |
|---|---|---|---|---|---|---|
| A | U | 0.00059 | 0.0 | 0.00047 | 0.00199 | 0.99999 |
|  | V | 0.00107 | 0.0 | 0.00085 | 0.00336 | 0.99983 |
| B | U | 0.04269 | 0.00182 | 0.02769 | 0.2899 | 0.99933 |
|  | V | 0.04933 | 0.00243 | 0.03188 | 0.29052 | 0.98641 |
| C | U | 0.0273 | 0.00075 | 0.02014 | 0.12429 | 0.99485 |
|  | V | 0.06088 | 0.00371 | 0.03971 | 0.32678 | 0.98903 |

Table 7.2: Model Performance for simulated 6-minute Online one-step forecasts for 72 hours of Pentland Simulation Data.

the HOFTS. The subsequent 3-day Online forecasts are displayed in Figure 7.5. It should be noted that the "spin-up" period for full online operation can be completed using only 15 days of reference data. Given most tidal energy sites will have ongoing data streams, this should not be prohibitive. It is clearly seen that despite increasing non-linearity between nodes $A \rightarrow C$, the model demonstrates a high level of agreement with the test data. Table 7.2 presents forecasting results. Forecasting accuracy decreases as the tidal signal is muddled by increased non-linearity and non-stationarity as expected.

Node A has minimal non-linearity and is easily modeled by HRA. The HA is able to account for 99.99% of the total signal variance, and the subsequent LRF and HOFTS are able to account for the remaining signal variance to an extremely high level of precision. There are increased levels of non-linearity and non-stationarity present at Node B resulting in steep fluctuations in current velocity as the tide direction changes. The HRA shows improved tracking ability through these sharp changes. While there is a tendency to over-predict the sharp changes in North-South velocity during the flood tide, the model is able to quickly return to agreement with the data. The increased accuracy around and after these turn points is reflected in the low RMSE values of .043 and .049 for the East-West and North-South currents, respectively.

Node C is characterized by extreme non-linear tidal flow interactions. As a result, we observe sharp changes in velocity during the slack-tide periods as well as

| Node | Prediction Data | RMSE | MSE | MAE | Max Error | $r^2$ |
|---|---|---|---|---|---|---|
| Cook Inlet | U | 0.06834 | 0.00467 | 0.05429 | 0.23987 | 0.99792 |
| | V | 0.0697 | 0.00486 | 0.0543 | 0.25361 | 0.84652 |
| Eastport | U | 0.02099 | 0.00044 | 0.01548 | 0.19906 | 0.95441 |
| | V | 0.03111 | 0.00097 | 0.02349 | 0.17139 | 0.99836 |

Table 7.3: Summary statistics for Eastport and Cook Inlet 6-minute predictions.

a shift in dominant velocity components from East-West to North-South. Despite these extreme characteristics, the proposed model is able to account for over 99.4% of the signal variance. While typical tidal energy sites are not characterized by such severe non-linearity, these results demonstrate the ability of the residual analysis method to capture non-linearities and have potential applications in archipelago regions where fast-moving tidal current mixing is prevalent.

### 7.5.2 Real Sites

In this section, we evaluate our model's performance on real data by simulating Online forecasts for two potential tidal energy sites in the United States. The Cook Inlet in Alaska is a high-interest tidal energy site that contains nearly a third of the nation's potential tidal energy resource (231). Data is taken from the East Foreland (ADCP Gauge ID: COI0503) from 18 May 2005, to 29 June 2005, at an approximate depth of $30.24m$, with measurements taken at a sampling rate of 6 minutes. The second set of data was taken from Estes Head (ADCP Gauge ID: EPT0003) located in Eastport, Maine. The Eastport data was collected from 01 July 2000, to 30 August 2000, with a sampling rate of 6 minutes.

#### 7.5.2.1 HRA Model Performance (Real Data)

As in the previous section, HA predictions were generated using 15 days of reference data. Linear Recurrent forecasts were run for seven days using only 36 hours (360 observations) as a training set size. The reduced training set size yielded increased performance for the sharp oscillations at the site at Estes Head. This was validated by the Hurst exponent which showed values less than 0.5 for training sets greater than 36 hours and a Hurst exponent of ≈ 0.6 for 36 hours. The performance statistics for the final 3-day combined forecasts are shown in Table 7.3.

The East Foreland is characterized by a dominant East-West current and a relatively slow-moving North-South current exhibiting a high degree of noise. It can be seen in Figure 7.6 that the HRA method demonstrates a high level of agreement for the East-West component through the sharp changes present during peak current velocities. This observation is confirmed by the $r^2$ value of .998. Due to the high level of noise, the North-South component exhibits less self-similarity which is confirmed by the reconstruction's Hurst exponents being less than 0.5. The adaptive construction of the HRA method is able to accommodate these signal characteristics by foregoing the LRF and predicting the "noise-free" residual using the HOFTS. It can be seen that the model prediction is suitable for a majority of test points and demonstrates reasonable agreement even with sharp velocity fluctuations.

The Eastport site is comparatively balanced, with both velocity components contaminated by severe non-linearity due to the site's location in a chain of archipelagos. As we demonstrated using the Pentland simulation, the HRA method is able to capture severe non-linear current interactions. As a result, the HRA method boasts Mean Average Errors of $0.015 m/s$ and $0.023 m/s$ for the East-West and North-South components respectively. The increased capacity of the the HRA method to capture these high-frequency variations relative to standard HA can be seen in Figure 7.7.

## 7.6 Power Predictive Performance

Accurate tidal current predictions 1 hour ahead are critical for the grid integration of tidal turbines. Typical acoustic doppler current profilers (ADCP) provide current measurements at a sampling rate of 6 minutes. As such, forecasting can either be carried out using multi-step or single-step forecasting. When we discuss time-series forecasting we typically only consider one-step forecasting; that is, forecasting a single time step ahead. Multi-step time series forecasting is a difficult task that

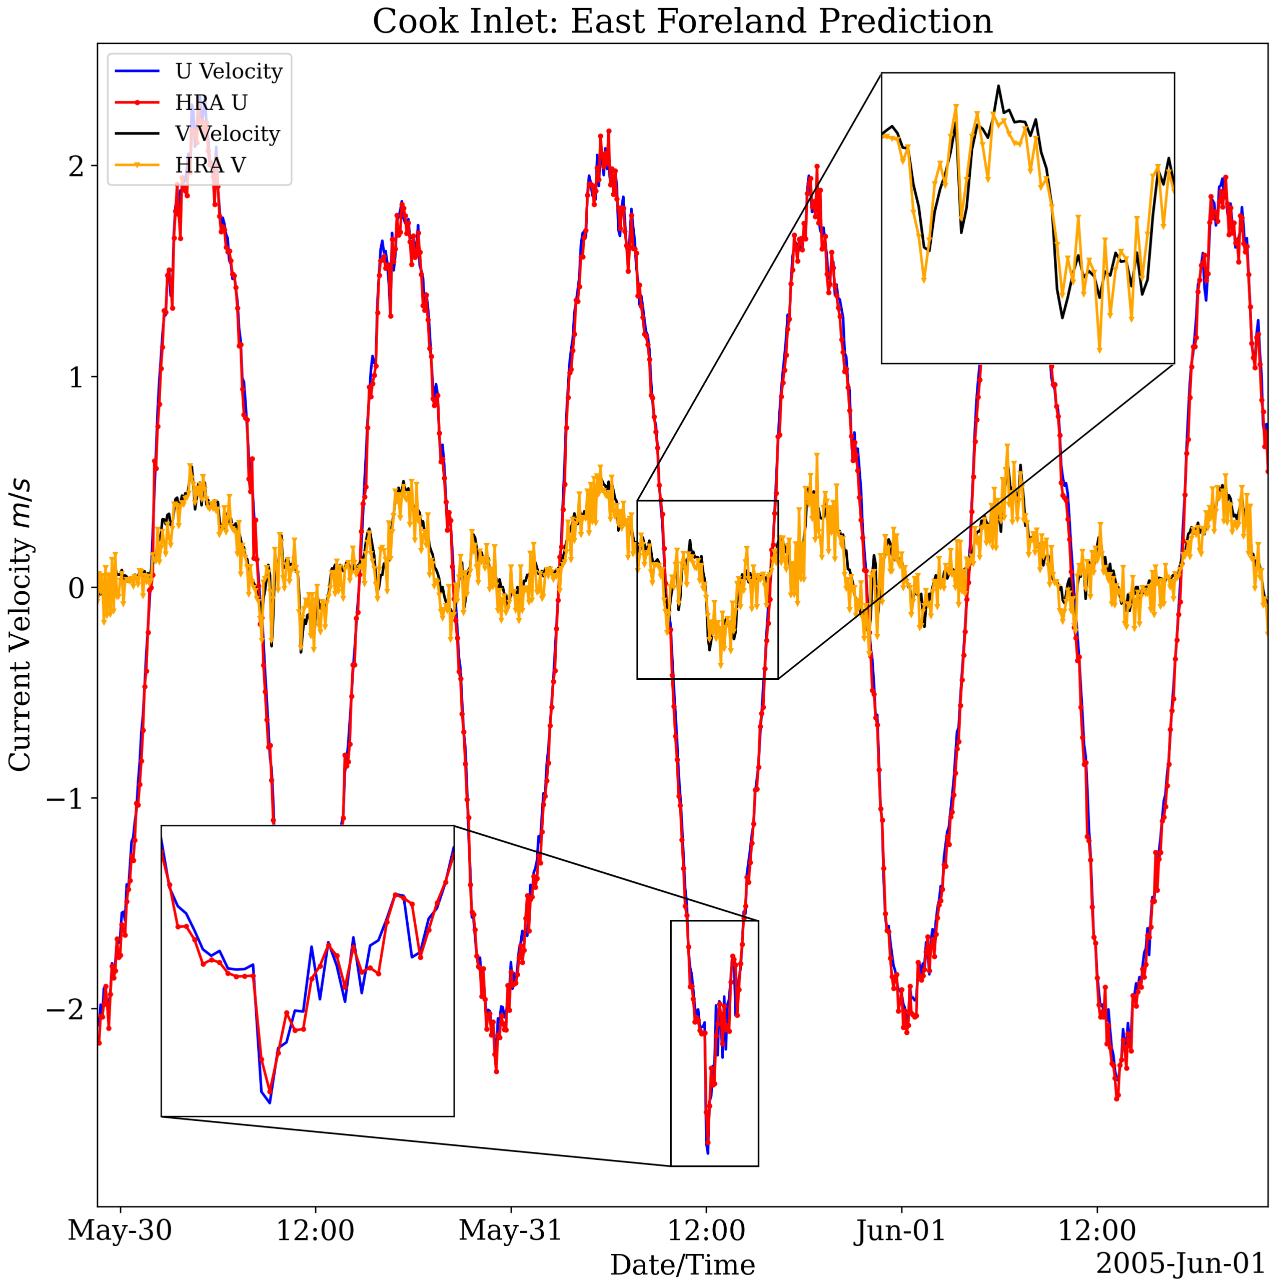

Figure 7.6: 72 hours of 6-minute online predictions for the ADCP at the East Foreland in the Cook Inlet, AK. Observed North-South and East-West velocities are shown in black and blue. HRA predictions for the North-South and East-West components are shown in orange and red.

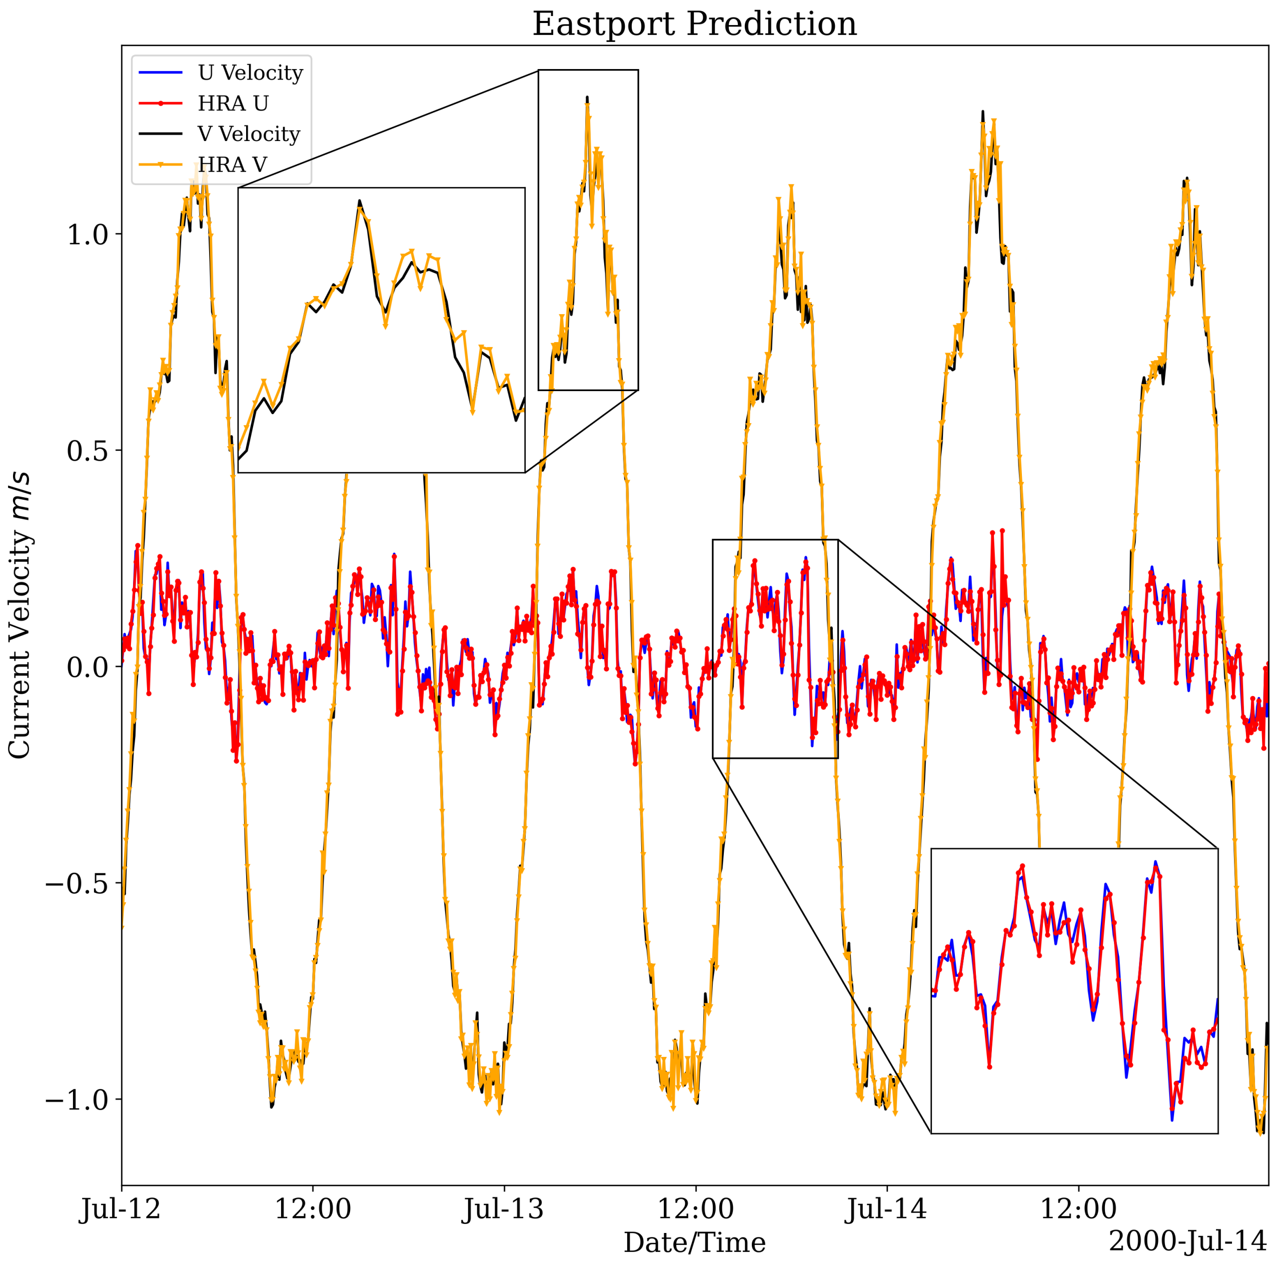

Figure 7.7: 72 hours of 6-minute online predictions for the ADCP at Estes Head in Eastport, ME. Observed North-South and East-West velocities are shown in black and blue. HRA predictions for the North-South and East-West components are shown in orange and red.

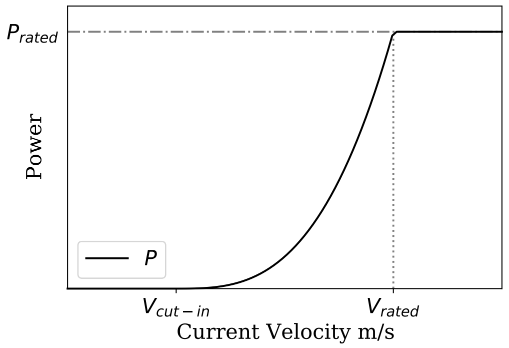

(a)

$$P = \begin{cases} 0 & 0 \le V_t < V_{cut-in} \\ 0.5C_p\rho A V_t^3 & V_{cut-in} \le V_t < V_{rated} \\ P_{rated} & V_t \ge V_{rated} \end{cases} \tag{7.12}$$

(b)

Figure 7.8: Panel (a) shows the theoretical power curve $P$ for a Tidal Current Turbine (TCT). Panel (b) shows the Power Equation where $C_p$ is the power capture coefficient $(0.4 - 0.5)$, $\rho$ is the water density, $A$ is the swept area of the blades $(81\pi)$, $V_t$ is the tidal current speed, $V_{cut-in}$ is the minimum current speed for turbine operation $(1m/s)$ and $V_{rated}$ is the rated operational speed $(3m/s)$ (8).

is performed using either a recursive or single-shot forecasting method (232). The recursive forecasting method makes multi-step predictions by recursively generating predictions one step ahead. As a result, recursive multi-step forecasting suffers from severe error propagation and rapid divergence. Single-shot multi-step forecasting predicts the entire multi-step horizon in one step. As such, the single-shot method is unable to consider the dependencies between observations (233). Due to these shortcomings and the relatively slow-changing nature of tidal currents, single-step 1-hour predictions are preferred for energy forecasts and are performed using down-sampled data.

It is important to understand the relationship between tidal currents and tidal energy generation in order to understand the importance of high-accuracy forecasts. Turbines are defined by a cut-in and rated velocity, $V_{cut-in}$ and $V_{rated}$. Below $V_{cut-in}$ the turbine is not operational and above $V_{rated}$ power output is constant. The power equation for a standard tidal turbine is given in Equation 12. A complete description of Tidal Power Generation systems is given in (8). Due to the cubic and discontinuous nature of the power equation, the forecasted power output is highly sensitive to errors in tidal current predictions. This cubic relationship is clearly illustrated in Figure 7.8 Panel (a). Despite high average accuracies, small errors in current velocity forecasts can result in significantly different power outputs. To

| Node | Prediction Data | RMSE | MSE | MAE | Max Error | $r^2$ |
|---|---|---|---|---|---|---|
| n4768 | U | 0.05287 | 0.00279 | 0.03388 | 0.23079 | 0.99953 |
|  | V | 0.06564 | 0.00431 | 0.0519 | 0.15373 | 0.98797 |
| n4496 | U | 0.09211 | 0.00848 | 0.07223 | 0.2879 | 0.98107 |
|  | V | 0.15409 | 0.02374 | 0.12071 | 0.38046 | 0.97955 |
| Cook Inlet | U | 0.05389 | 0.0029 | 0.03706 | 0.25675 | 0.99885 |
|  | V | 0.02507 | 0.00063 | 0.01677 | 0.08573 | 0.97461 |
| Eastport | U | 0.03073 | 0.00094 | 0.02083 | 0.12205 | 0.85397 |
|  | V | 0.05301 | 0.00281 | 0.04074 | 0.22691 | 0.99547 |

Table 7.4: Pentland, Eastport, and Cook Inlet 1-Hour performance results for 15-days of simulated Online forecasts.

illustrate our model's power predictive performance results for 1-hour tidal current predictions and the associated power forecasts are provided.

#### 7.6.0.1 HRA Model Performance (1-Hour Horizon)

The effectiveness of HRA is evaluated for 1-hour predictions using downsampled data from the Pentland Firth simulation, Eastport, and the Cook Inlet. HA predictions are generated using 30 days of reference data from each site (720 observations). SSA-LRF forecasts are run for 15 days on the HA residuals using a training set length of 15 days (360 observations) and a window length of 150. The window length was determined based on the recommendations given in (234) and validated using the Hurst exponent. The SSA reconstructed high-PE signal is presented to the HOFTS and the final prediction is aggregated. Due to the smoothing nature of downsampling, the behavior of the resultant time series and computed HA residuals is different from that of the 6-minute data. For the 6-minute data, the reconstructed SSA low-PE signal's Hurst exponents were all found to be suitable for the LRF forecasts with the exception of the Cook Inlet's North-South component. Conversely, for the downsampled 1-hour data, all reconstructed signals were suitable for LRF with the exception of the Eastport East-West component. As such, after removing true noise components with $PE > 0.99$ the combined low and high-PE signals for the Eastport East-West component were predicted by the HOFTS. Table 7.4 presents the performance indicators for the 15 days of online forecasts.

It can be clearly seen in Table 7.4 that the HRA method yields similar performance for 1-hour predictions as it did for 6-minute predictions. Furthermore, the model demonstrates high performance across sites exhibiting significantly different characteristics and degrees of non-linearity. The high performance across assessment metrics over a full Spring-Neap cycle verifies that the HRA method is suitable for 1-hour forecasts.

Due to the low self-similarity in the SSA reconstructed signal for the Eastport East-West component, the increased reliance on the HOFTS prediction results in a decreased prediction accuracy. Despite this, the prediction retains a low MAE of 0.02 and a maximum error of 0.122. Hence, we have shown that by combining the stability of numerical methods with the flexibility of ML methods, HRA provides increased reliability in the presence of unknown conditions and results in a reduction of maximum error compared to conventional ML models.

Power predictions are computed using Equation 7.8b. Figure 7.9 shows a 10-day subset of the 15-day forecast as well as the associated Absolute Error distributions comparing the HRA method with traditional HA. It is clear that HRA can significantly reduce both the maximum and mean prediction error for real-time power forecasting when compared to numerical schemes. As such, we present the HRA method as a reliable tool to augment the numerical schemes currently employed by tidal energy practitioners.

## 7.7 Conclusion

To the best of our knowledge, this chapter presents the first online tidal current prediction method. The hybrid ML approach combines the reliability of classical HA with the improved localization and flexibility of ML methods. The primary developments from this chapter are as follows:

1. The experimental results verify the HRA method's viability for both 6-minute and 1-hour online tidal current forecasting. The method is robust to non-

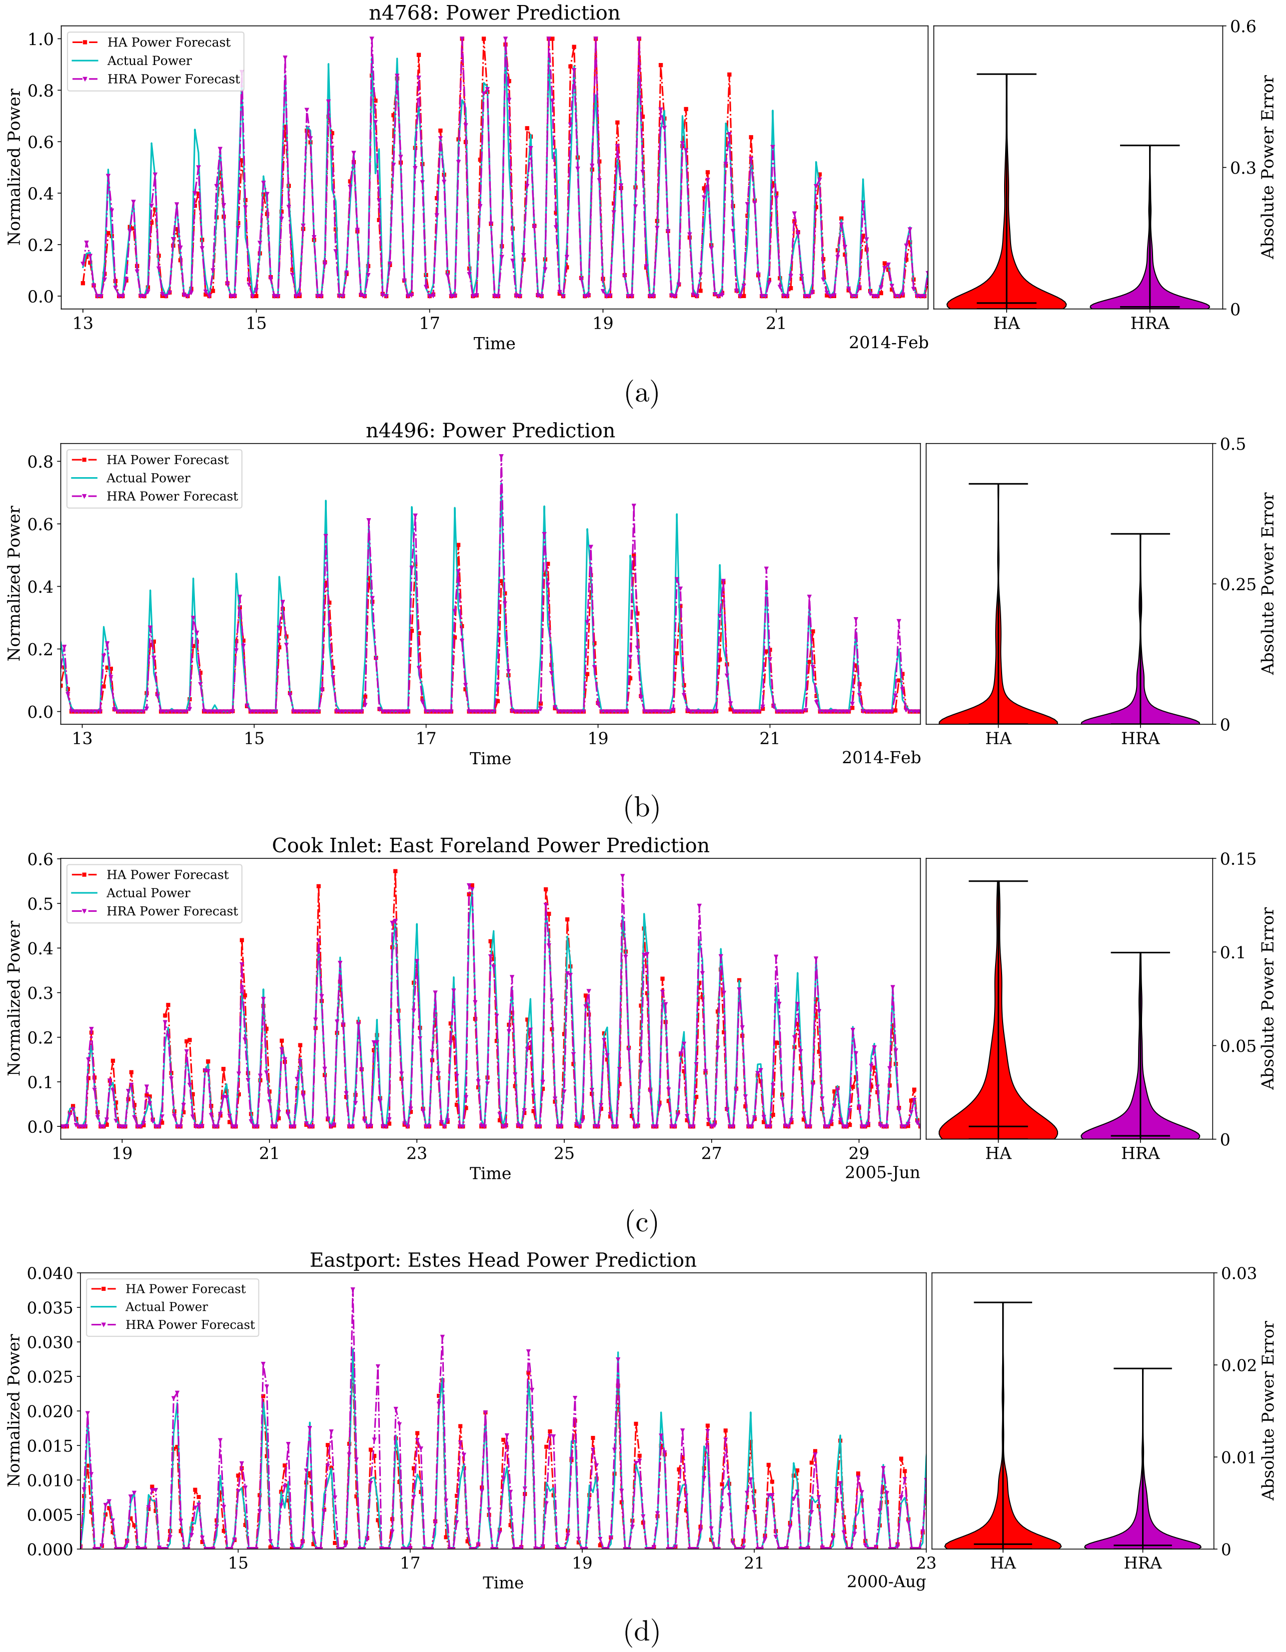

Figure 7.9: Power prediction comparisons between Harmonic Analysis and the HRA method for 1 Hour Downsampled data over a Spring-Neap Cycle. (a) Pentland n4768 (b) Pentland n4496 (c) Cook Inlet, East Foreland (d) Eastport, Estes Head. The corresponding violin plot shows the distribution of absolute error values with whiskers indicating the mean and maximum error values. Power is normalized such that 1 corresponds to currents above $V_{rated}$ and 0 below $V_{cut-in}$.

linearity and non-stationarity and has been shown to work across a diverse set of sites.

2. This is the first time an automated component selection criteria has been used for SSA-LRF. The unsupervised binary grouping approach, augmented by the moving Hurst Exponent, results in improved LRF accuracy and reliability. The method shows promise as a general tool for SSA-LRF component selection; however, further study is necessary.
3. We have demonstrated that the HRA approach can achieve better accuracy over traditional HA and can therefore be used to augment the forecasting schemes at existing tidal energy sites. Moreover, the proposed model realizes accurate online predictions for both tidal currents and the resultant power generation. Comparison with state-of-the-art online prediction methods remains to be determined.

Future work will look to expand the HRA method to incorporate multiple data streams, including adjacent ADCPs and meteorological forcing variables (wind stress, atmospheric pressure, etc.). The inclusion of these parameters can enhance the model's predictive capabilities and increase tolerance to severe weather. Data for the three nodes (n500, n4768, and n4496) from the Pentland simulation will be made available and can be used to study the behavior of fast-moving tidal currents as well as act as a standard benchmark for future models.

## Additional Information

### 7.7.1 Performance Measures

Root Mean Squared Error (RMSE):

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{7.13}$$

Mean Squared Error (MSE):

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{7.14}$$

Mean Absolute Error (MAE):

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{7.15}$$

Maximum Absolute Error:

$$Max\ Error = \max(|y_i - \hat{y}_i|) \tag{7.16}$$

R-Squared ($r^2$):

$$r^2 = 1 - \frac{\frac{1}{n}\sum_{i=1}^{m}(y_i - \hat{y}_i)^2}{\frac{1}{n}\sum_{i=1}^{m}(y_i - \bar{y}_i)^2} \tag{7.17}$$

Note $n$ is the number of points sampled, $y_i$ is the actual value, $\hat{y}_i$ is the predicted value, and $\bar{y}_i$ is the mean value.

### 7.7.2 Visual Summary of Permutation Entropy

We present herein a visual explanation of Permutation Entropy in Figure 7.10. Panel (a) illustrates the relationship between the order $n$ and time lag $\tau$ with the timestep shown on the x-axis and the corresponding unitless observation values shown on the y-axis. The importance of $\tau$ in determining the frequency characteristics analyzed can be clearly seen. Panel (b) is the corresponding normalized histogram of permutation occurrences for (a). From this, it can be seen how the partitioned segments are classified.

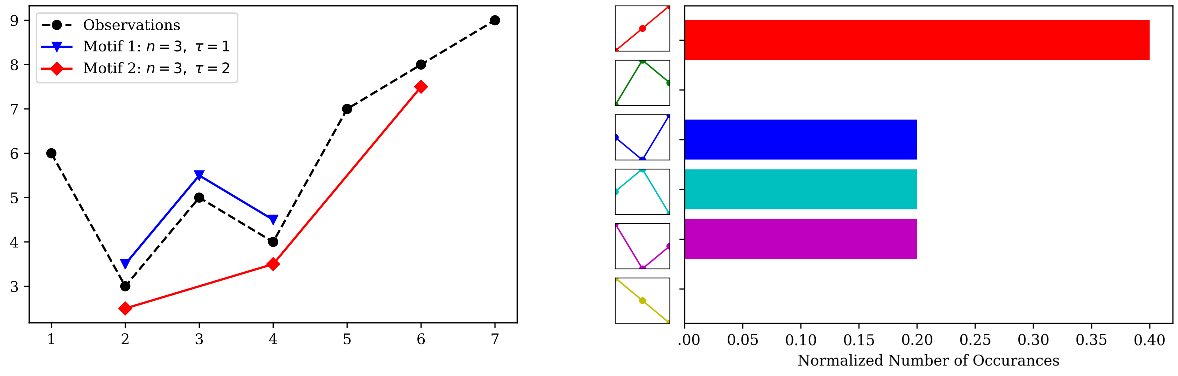

(a) (b)

Figure 7.10: Simplified example of the relationship between $n$ and $\tau$ for PE.

| Node | Number | Latitude | Longitude |
|---|---|---|---|
| A | 500 | $59.8599068915N^{\circ}$ | $4.9797111830W^{\circ}$ |
| B | 4768 | $58.6557103353N^{\circ}$ | $3.1391667722W^{\circ}$ |
| C | 4496 | $58.6918591258N^{\circ}$ | $3.1296765346W^{\circ}$ |

Table 7.5: Location of Case Study Nodes at Pentland Firth.

### 7.7.3 Node Locations and Simulation Designations

### 7.7.4 Machine Learning Benchmarks

#### 7.7.4.1 ML Hyperparameters and Architecture

An overview of the model architecture and hyperparameters for each ML algorithm used for benchmarking is provided below.

- **RF:** RF regression was implemented using the sci-kit-learn RandomForestRegressor module (235). Predicting based on the previous 8 values yielded the best performance in our testing.
- **SVM:** SVM was implemented using the sci-kit-learn SVM module. A radial basis function (RBF) kernel was selected based on the recommendations from (197). The SVM produced the highest performance when training to predict off of the 4 previous values.
- **ANN:** The backpropagation ANN was implemented using Keras (236). The model used an input layer with 8 neurons and two hidden layers with sizes 16 and 8 respectively. The ADAM optimizer was used with a learning rate of

| Site | Method | Prediction Data | RMSE | MSE | MAE | Max Error | r$^2$ |
|---|---|---|---|---|---|---|---|
| Node A | UTide | U | 0.0412 | 0.0017 | 0.0364 | 0.0749 | 0.9494 |
| | | V | 0.035 | 0.0012 | 0.0321 | 0.0544 | 0.7107 |
| Node A | RF | U | 0.0059 | **0.0**** | 0.0036 | 0.0321 | 0.9992 |
| | | V | 0.0012 | **0.0**** | 0.0009 | 0.0043 | **0.9998**** |
| Node A | SVR | U | 0.069 | 0.0048 | 0.0581 | 0.1411 | 0.7933 |
| | | V | 0.0618 | 0.0038 | 0.0558 | 0.0983 | -5.9303 |
| Node A | ANN | U | 0.0012 | **0.0**** | 0.0008 | 0.0051 | **1.0*** |
| | | V | **0.0011*** | **0.0**** | 0.0009 | 0.0046 | **0.9998**** |
| Node A | LSTM | U | 0.0036 | **0.0**** | 0.0021 | 0.0207 | 0.9997 |
| | | V | 0.0021 | **0.0**** | 0.0019 | 0.0054 | 0.9994 |
| Node A | HRA | U | **0.0006** | **0.0**** | **0.0005** | **0.002** | **1.0*** |
| | | V | **0.0011*** | **0.0**** | **0.0008** | **0.0034** | **0.9998**** |
| Node B | UTide | U | 0.2118 | 0.0449 | 0.1734 | 0.7643 | 0.8801 |
| | | V | 0.3725 | 0.1388 | 0.3181 | 0.957 | 0.839 |
| Node B | RF | U | 0.0861 | 0.0074 | 0.0542 | 0.4807 | 0.9836 |
| | | V | 0.1627 | 0.0265 | 0.0974 | 0.7945 | 0.9779 |
| Node B | SVR | U | 0.0822 | 0.0068 | 0.065 | 0.4287 | 0.9846 |
| | | V | 0.1676 | 0.0281 | 0.1115 | 0.8274 | 0.9757 |
| Node B | ANN | U | 0.0685 | 0.0047 | 0.042 | 0.3891 | 0.9897 |
| | | V | 0.2393 | 0.0573 | 0.1466 | 1.0067 | 0.946 |
| Node B | LSTM | U | 0.0722 | 0.0052 | 0.045 | 0.5516 | 0.9886 |
| | | V | 0.1517 | 0.023 | 0.0821 | 0.7839 | 0.9806 |
| Node B | HRA | U | **0.0427** | **0.0018** | **0.0277** | **0.2899** | **0.9993** |
| | | V | **0.0493** | **0.0024** | **0.0319** | **0.2905** | **0.9864** |
| Node C | UTide | U | 0.3942 | 0.1554 | 0.33 | 1.076 | 0.9676 |
| | | V | 0.1473 | 0.0217 | 0.1123 | 0.4348 | 0.9286 |
| Node C | RF | U | 0.1293 | 0.0167 | 0.1013 | 0.52 | 0.995 |
| | | V | **0.0497** | **0.0025** | **0.0318** | 0.3411 | 0.9888 |
| Node C | SVR | U | 0.0914 | 0.0084 | 0.0725 | 0.4204 | **0.9975** |
| | | V | 0.0563 | 0.0032 | 0.0483 | 0.3649 | 0.9871 |
| Node C | ANN | U | 0.13 | 0.0169 | 0.0772 | 0.7508 | 0.9949 |
| | | V | 0.0724 | 0.0052 | 0.0427 | 0.6064 | 0.9865 |
| Node C | LSTM | U | 0.2049 | 0.042 | 0.1563 | 0.6391 | 0.988 |
| | | V | 0.0785 | 0.0062 | 0.0568 | 0.3549 | 0.9753 |
| Node C | HRA | U | **0.0273** | **0.0008** | **0.0201** | **0.1243** | 0.9949 |
| | | V | 0.0609 | 0.0037 | 0.0397 | **0.3268** | **0.989** |

Table 7.6: Machine Learning comparisons for Pentland Firth Simulation sites. Model predictions are evaluated over 3 days on data sampled at 6-minute intervals. Note, the best values are highlighted in bold for each performance measure and current direction respectively. * denotes values for which two models are tied for best performance, and ** denotes values for which multiple values are tied for best performance.

0.0001 and was trained for 500 epochs. We found that the best performance resulted from training the model to predict based off of the 8 previous values.

- **LSTM:** The LSTM was implemented using the Keras LSTM module. The

| Site | Method | Prediction Data | RMSE | MSE | MAE | Max Error | $r^2$ |
|---|---|---|---|---|---|---|---|
| Eastport | UTide | U | 0.0741 | 0.0055 | 0.0588 | 0.2527 | -0.2413 |
| | | V | 0.3003 | 0.0902 | 0.255 | 0.7226 | 0.8621 |
| Eastport | RF | U | 0.0479 | 0.0023 | 0.0369 | 0.1904 | 0.6622 |
| | | V | 0.057 | 0.0033 | 0.0437 | 0.2026 | 0.9939 |
| Eastport | SVR | U | 0.0468 | 0.0022 | 0.0366 | 0.1897 | 0.6429 |
| | | V | 0.0537 | 0.0029 | 0.0412 | 0.2075 | 0.9946 |
| Eastport | ANN | U | 0.046 | 0.0021 | 0.0351 | 0.1819 | 0.7024 |
| | | V | 0.5586 | 0.312 | 0.4759 | 1.358 | -4.9505 |
| Eastport | LSTM | U | 0.0462 | 0.0021 | 0.0356 | **0.1761** | 0.6942 |
| | | V | 0.0549 | 0.003 | 0.0425 | 0.2115 | 0.9945 |
| Eastport | HRA | U | **0.021** | **0.0004** | **0.0155** | 0.1991 | **0.9544** |
| | | V | **0.0311** | **0.001** | **0.0235** | **0.1714** | **0.9984** |
| Cook Inlet | UTide | U | 0.4408 | 0.1943 | 0.3901 | 1.0072 | 0.899 |
| | | V | 0.1265 | 0.016 | 0.0995 | 0.3494 | -0.1896 |
| Cook Inlet | RF | U | 0.0867 | 0.0075 | 0.065 | 0.3339 | 0.997 |
| | | V | 0.074 | 0.0055 | 0.0555 | 0.2807 | 0.7846 |
| Cook Inlet | SVR | U | 0.0857 | 0.0073 | 0.065 | 0.3258 | 0.9971 |
| | | V | 0.0842 | 0.0071 | 0.0635 | 0.305 | 0.6719 |
| Cook Inlet | ANN | U | 0.0915 | 0.0084 | 0.0668 | 0.3988 | 0.9967 |
| | | V | 0.1044 | 0.0109 | 0.0805 | 0.3203 | 0.785 |
| Cook Inlet | LSTM | U | 0.0961 | 0.0092 | 0.0747 | 0.3806 | 0.9963 |
| | | V | 0.0746 | 0.0056 | 0.056 | 0.3105 | 0.7728 |
| Cook Inlet | HRA | U | **0.0683** | **0.0047** | **0.0543** | **0.2399** | **0.9979** |
| | | V | **0.0697** | **0.0049** | **0.0543** | **0.2536** | **0.8465** |

Table 7.7: Machine Learning comparisons for real sites. Model predictions are evaluated over 3 days on data sampled at 6-minute intervals. Note, the best values are highlighted in bold for each performance measure and current direction respectively.

model used an input layer with 8 neurons and one hidden layer with size 8. The ADAM optimizer was used with a learning rate of 0.00001 and was trained for 500 epochs. We found that the best performance resulted from training the model to predict based on the 8 previous values.

It should be noted that training times for the LSTM and ANN models exceeded 10 minutes and can therefore not be used in an Online manner. In our testing, we found that training separate models to predict the U and V components respectively yielded superior results for all models.

### 7.7.5 Model Parameters

The SSA parameters used for each site and sampling rate are reported in Table 7.8. It should be noted that grid partitioning is used for the HOFTS on all sites and is set to 15 partitions.

| Site | Sampling Rate | Train Set Length (samples) | Window Length (samples) |
|---|---|---|---|
| Pentland Firth | 6 Minute | 720 | 350 |
| | 1 Hour | 360 | 150 |
| Cook Inlet | 6 Minute | 360 | 150 |
| | 1 Hour | 360 | 150 |
| Eastport | 6 Minute | 360 | 150 |
| | 1 Hour | 360 | 150 |

Table 7.8: Model parameters for SSA decomposition.

# 8 Conclusion

The purpose of this thesis was to enhance the two main empirical methods of tidal analysis and prediction through the introduction of data-driven machine learning. As outlined in the introduction, the limitations of the conventional signal processing methods employed for these tasks have reduced their effectiveness, and for the response method, inhibited its use altogether. In this final chapter, I explicitly define the gaps filled by this research and discuss the limitations and new avenues of research brought by these methods.

## 8.1 Summary of studies

This thesis first considers the challenge of performing harmonic analysis on sparsely sampled and strongly aliased satellite altimetry data from the new SWOT mission. The design of the method in Chapter 3 was tailored to maximize the scientific return from the CAL/VAL orbit which lasted just over 100 days. Exploitation of such sparse data is challenging due to the combination of the temporal sparsity and the noise present in the data. Chapter 3 introduces three major innovations which improve the ability to use these data. First, I develop a variational Bayesian estimator to replace conventional least-squares estimation. This estimator is more robust to various types of noise, and explicitly models uncertainty. Using this uncertainty information, I develop a Bayesian pruning procedure which identifies and removes data-points which are statistically not tidal. I further develop a spatially coherent inductive bias which allows the model to exploit the information from adjacent locations. These three innovations work together to improve the robustness of the estimation procedure to noise, and demonstrate significant improvements in derived tides over both conventional least-squares and state-of-the-art assimilative models. Chapter 4 demonstrates why this approach is critical to studying non-tidal processes. The strength of conclusions in that study were only made possible by the methodological developments in Chapter 3.

Second, no single method exists for the analysis and prediction of arbitrary tidal processes. That is, processes which are driven by tides, but are not necessarily stationary due to modulation from some non-tidal forcing. While Munk and Cartwright's response method could in principle be applied to such situations, the limitations of the lumped and sequential analysis procedures made this largely intractable. Chapter 5 develops an automated analysis procedure within the classical response method avoiding these limitations altogether. In so doing, the method can be applied to analyze and predict any tidal process if the time varying forcing is provided to the model[1]. A key finding in Chapter 5 is that by accounting for this forcing, a superior estimate of the tide is obtained. I further develop this approach for application to tidal currents in Chapter 6. Tidal currents are often more complex than tidal heights, and reference series are often comparatively shorter. Both of these characteristics are more naturally accounted for within a response framework as described. The major innovation in Chapter 6 is the development of a coupled response model, whereby the orthogonal tidal current velocity components are estimated simultaneously using a set of partially shared weights. I demonstrate how this significantly improves the accuracy of estimated tides over harmonic analysis for shortened reference series. Further, even for longer reference series, the coupled response model demonstrates a reduction in absolute error of approximately 10% over harmonic methods. Hence, for engineering applications where accuracy is paramount, this work demonstrates the coupled response approach should be favored over standard harmonic methods. I also demonstrate how the response approach can be used to improve the accuracy of the derived harmonic constituents from short reference series. In this way, one need not compromise on the familiar interpretability brought by the harmonic method when conducting a response analysis.

Lastly, Chapter 7 addresses the problem of short-term tidal-energy forecasting. In contrast to long-term forecasting (performed using the methods above), short-term forecasting is made difficult by transient, and stochastic time-dependent processes

---

[1] Experiments (not included in this thesis) have shown that it is possible to still perform analyses on non-stationary tidal processes even if the external forcing cannot be specified. All one needs to do is to add in the normalized time (from 0-1) as an additional input function. Prediction is naturally limited to times immediately after the training period.

which are not necessarily feasible to capture within harmonic or response frameworks. This is important as these processes can significantly alter the yield of tidal energy devices. The primary contribution in this work is the development of an 'Online' training procedure which allows the model to respond to recent changes in conditions. By pairing this with a harmonic analysis, the hybrid method can account for both the time-invariant (harmonic) and time-varying (stochastic) components of the observed currents. In so doing, the model can significantly improve the estimated tidal currents and subsequent power production of standard methods. An advantage of this approach is that the residual analysis can be conducted on any model, and is thus applicable to both numerical and response models as well. This is critical as it can reduce the errors produced by these models which assume time-invariance if a transient process is present.

## 8.2 Contributions, limitations, and implications for future work

In addition to the contributions above, this work opens doors for several new avenues of research which I outline below. While Chapters 3 and 4 considered the application of the VBayes estimator for satellite altimetry, the approach has considerable utility across other applications.

A key advancement in Chapter 3 is the so-called Bayesian pruning procedure. This approach leverages the tidal estimator's probabilistic construction to identify measurements which significantly deviate from the estimated tidal signal. This provides a statistically-grounded outlier removal procedure as called for in prior work (33). As noted in Chapter 2, even when analyzing in-situ gauges, large errors (relative to the true tidal signal) occur due to bio-fouling and sensor dropout. While these errors are large with regard to the tidal signal, they rarely exceed the statistical thresholds necessary to justify their removal relative to the overall distribution of values. By evaluating the statistical relevance of these data relative to the tidal signal, removal of these values is made trivial. Using this approach, I am in the process of reprocessing the GESLA tide gauge dataset to provide an additional flag for outliers missed using their standard processing (51).

Even when data is abundant, spurious data can still severely impact the accuracy of a tidal analysis from long standing gauges (36). These errors can impact the accuracy of global tide model comparisons, and create further problems if assimilated into numerical models directly. The pruning procedure has the potential to partially mitigate these errors by pre-filtering data prior to harmonic analysis (either using conventional or Bayesian estimators). Future work will look to integrate this approach into the processing pipeline of major tidal constituent datasets such as TICON-3 (52).

The pruning procedure illustrates one way in which the improved uncertainty estimation from the VBayes estimator can be utilized. A second, and important area, is in producing accurate uncertainty estimates for global tide models. In this vein, replacement of conventional least-squares estimation with the VBayes approach has the potential to immediately make progress towards this objective. As noted in Chapter 3, the assumption of a Gaussian noise distribution does not hold in all cases. This assumption was necessary for the extremely sparse SWOT data; however, when data is abundant (e.g. longstanding in-situ gauges), modification of the approach to handle red-noise represents an important area of future work. A simple modification is to modify our approach to instead model the noise as a mixture of Gaussian distributions (237).

The usage of ordinary least-squares is ubiquitous across nearly all geophysical disciplines. As discussed, the assumptions made by OLS estimators can lead to severe problems depending on the nature of the data. The VBayes approach developed in Chapter 3 and applied in Chapter 4 provides a principled approach to handling conditions present in real-data and has conceptual advantages over typical regularization methods. Hence, future work should look to evaluate whether methods which currently rely on least-squares solutions could benefit from integration of VBayes estimators.

Central to this work has been developing approaches which enhance our ability to study tidal processes. The case-studies in Chapter 5 illustrate a few ways in which the automated response method can be useful, however, it remains to be said where else this should be used. The primary advantage of the response approach is the ability to explicitly model how tides are modulated by external forcing. I envision this being useful for studies into tidal change both as it relates to sea-level rise and stratification. Globally, spatially coherent trends in tides are being observed with links to both of these processes (and more) (238; 239; 240). Rather than estimating the trend in tidal constituents from these observations, one could instead attempt to learn the time-invariant response to various exogenous inputs (e.g. stratification and/or mean-sea-level). Initial testing (not in this thesis) has shown that including the mean-sea level as an additional input can yield improvements in model predictive skill greater than if one simply de-trends the data. Hence, further investigation of this, and the coupled response to both changing sea-levels and stratification is needed.

As pointed out by Cartwright in his initial work on surges (16), the response method could be utilized for operational surge forecasting assuming (i) the non-linear response to the meteorological forcing could effectively be learned, and (ii) the forcing could be accurately forecasted. While the latter problem has been solved for decades, this thesis is the first to address the first problem. Work not included in this thesis has demonstrated that numerical surge models can be augmented using my response approach by feeding the numerical model in as an additional input function. Under this construction the model learns the coupled tide-surge response which is poorly resolved in these models. Initial independent testing by the UK National Oceanography Centre and Environment Agency has demonstrated reductions in mean absolute error over their operational forecast system by more than

70% at Sheerness, UK. Independent testing by operators of the Dutch surge barriers have shown similar improvements using RTide, and are currently implementing the model operationally. While these applications are exciting, I believe the response approach can have even greater impact in generating operational forecasts in regions which lack sophistical numerical forecast systems. Initial testing, building off the meteorological forcing case study in Chapter 5, has shown that by incorporating meteorological forcing over a larger spatial region, pure response based models can produce comparable or better performance to conventional numerical forecasts. If such an approach could be combined with satellite altimetry data, models could be developed anywhere in the world without requiring accurate (or any) bathymetry data.

Current altimetry missions such as Sentinel-1, SWOT, and in the future ODYSEA, are/will be providing high-accuracy SAR based Doppler shift sea-surface measurements. Recent studies have shown that these measurements can actually be analyzed using standard harmonic analysis to estimate the tidal currents (241). The sparsity of satellite repeat measurements makes this task particularly difficult. The ability to use less data to estimate currents using the coupled response method in Chapter 6 could prove promising for application to these data and overcoming the unfavorable sampling from these satellites respective orbits.

The basic Volterra network proposed in Chapter 5 is grounded by its connection to the Volterra series; however, it is not the only model which can be employed for weight estimation. In fact, essentially any ML model can be substituted in depending on the objective. Thus, it is worth asking what the limitations of the Volterra networks are to understand where future work should focus. To me, the greatest limitation of the Volterra network is its lack of uncertainty quantification – something which remains a limitation of most deep-learning methods. To overcome this, a family of methods called ‘neural processes’ could prove useful (242). In this framework, the response function (and its uncertainty) is explicitly modeled as a distribution of functions rather than a deterministic mapping as is currently used in RTide. In so doing, one can quantify both the aleatoric and epistemic uncertainty in the derived response function which yields well calibrated uncertainties over the model prediction. This method has shown promise on initial testing in the above mentioned surge forecast applications.

With these avenues of future work acknowledged, it is worth considering where and when the RTide approach developed in Chapters 5 and 6 should be used. The fundamental assumption made by the response method is that the system under study is time-invariant. If time varying processes modulate this response then this assumption is invalid. Examples of this would be if one were to conduct a purely gravitational response analysis in a tidal river. Chapter 6 showed that when this time-invariance is invalidated, the response approach generally performs worse that conventional harmonic analysis. The great advantage of the RTide approach, however, is the ability to directly incorporate such forcing in the analysis. Appropriate usage of RTide therefore relies on users recognizing what physical processes are modulating the tidal response, and being able to provide time-series of these processes as additional forcing. The approximate tidal to long-period energy ratio advocated in Chapter 6 provides a starting point for recognizing when these modifications are necessary. Another limitation of the RTide approach is that it deliberately removes the ability to explicitly define the interactions between parameters *a priori*. As shown, this was necessary for automating the response method; however, it limits the ability to control which interactions (linear or nonlinear) the model learns. Such a feature may be useful for those conducting fundamental tidal research. This thesis develops several *ex post facto* approaches which can be used to explore the interactions of learned models but provides no solution to the former problem aside from the specification of input functions. Hence, if this control is of interest to users then the classical response method may be more appropriate.

## 8.3 Final thoughts

This work sets itself apart from a majority of machine-learning approaches by looking to augment rather than replace the conventional forms of tidal analysis and prediction. The central idea this work advances is that by understanding the assumptions our models make about the world, we can construct methods which can mitigate and reduce their over simplistic assumptions, and even exploit these assumptions to make better use of the available data. As shown, this yields models which are more resilient to the inevitable noise present in the real world and consequently more performant. The shear breadth of applications considered in this work is evidence of the value these approaches have.

However, as outlined in the previous discussion, the work is not done. There are countless ways in which tidal processes are changing as a consequence of climate change. These changes have major implications for coastal communities and industries operating in these environments. It is my hope that the tools developed in this thesis, and published as the open-source RTide and VTide Python packages, can help to partially address these challenges. With new observational technologies coming online in the coming years and exponentially more data now available, it is necessary for new methods, based on these classical approaches, to be developed. The philosophy advocated in this thesis provides a roadmap for accomplishing this through the smart infusion of physics-based inductive biases with data-driven machine learning.

# A Supporting Information for Tidal corrections from and for SWOT using a spatially coherent variational Bayesian harmonic analysis

## Contents of this appendix

1. Text A1
2. Figures A1-A6
3. Table A1

The following supplementary materials provide additional information and figures to support the claims made in Chapter 3 along with some interesting results. First, we show the results of the evolutionary optimization when optimized for the constituents (M2, N2, O1, and M4) in Figure A.1. Next, model performance is evaluated when varying the sampling around the 35 days sampling period of the Envisat satellite in Figure C.2. The result for the resolution of the S2 constituent is quite striking. It can be seen that the OLS estimator suffers from the expected “aliasing” with an RMS error of $10^{12}$. This can be seen to degrade the quality of the derived K1 constituent also. Conversely, the VBayes RMS error for S2 when averaged over the entire simulated TICON-3 dataset is less than 9cm. Similar robustness for the additional constituents is seen also. We believe this is an important result as the inclusion of Envisat measurements in empirical analysis is known to degrade the quality of S2 and other constituents. Figure A.3 compares the performance of the standard VBayes, OLS, along with the Spatial VBayes and Spatial OLS under 100 dB WGN. Figures A.4 and A.5 compare the performance of the empirical approaches and FES on the 22 gauges contained within passes 003 and 016 when performing direct empirical analysis (SLA + Tide Correction). Finally figure A.6, visualizes the spatially coherent inductive bias.

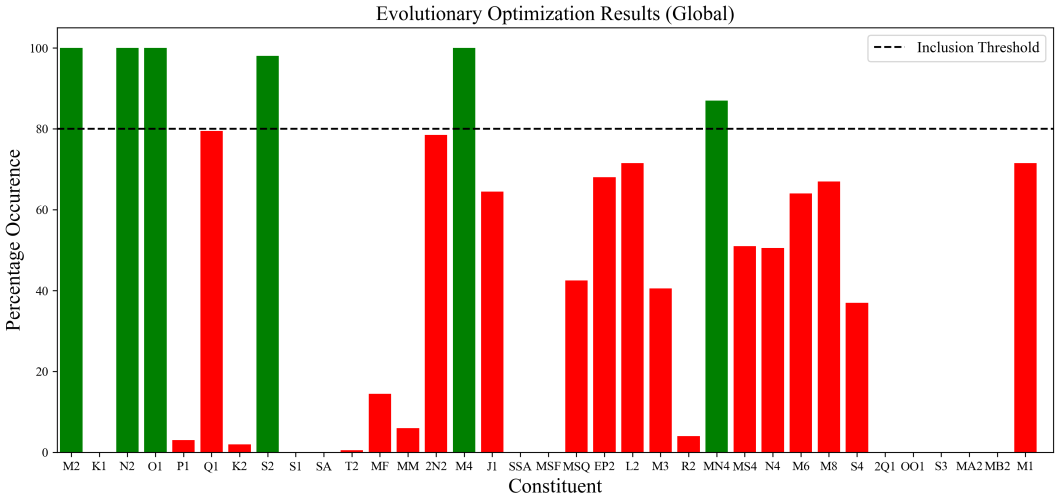

Figure A.1: **Percentage occurrence of different constituents from the top 20 evolutionary optimization solutions**. The final solution is determined by the frequency of occurrence when optimized for both the relative percentage error and RMS error and the primary constituents (M2, N2, O1, and M4). The 80% threshold is arbitrarily defined based off of empirical analysis of the spatial coherence when applied to real SWOT data.

| **Gauge 1** | **Gauge 2** |
|---|---|
| (51.55, 357.014) | (51.55, 357.013) |
| (50.608, 357.553) | (50.608, 357.552) |
| (51.508, 357.287) | (51.5, 357.272) |
| (56.623, 353.937) | (56.623, 353.936) |
| (54.651, 356.433) | (54.65, 356.433) |
| (54.843, 354.88) | (54.842, 354.881) |
| (51.215, 356.865) | (51.217, 356.867) |

Table A.1: List of duplicate TICON-3 gauges contained with passes 003 and 016.

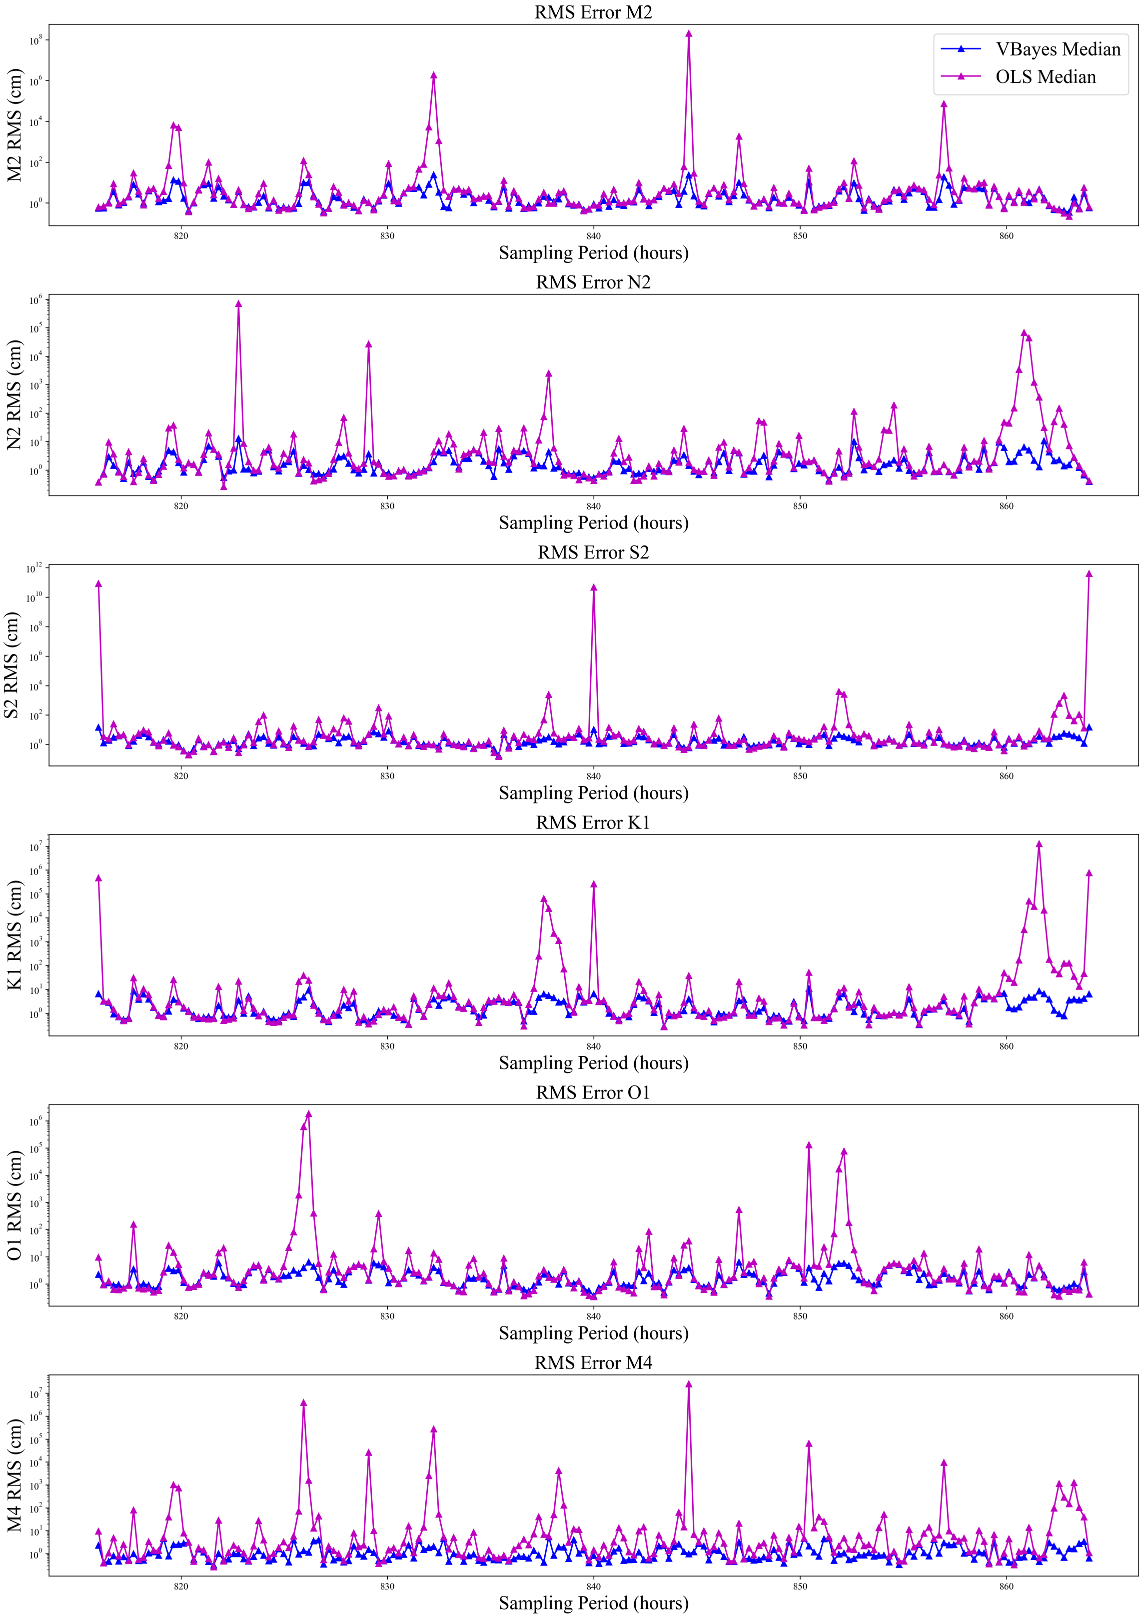

Figure A.2: **Comparisons of median RMS (cm) for OLS and VBayes for several constituents when varying sampling period around Envisat period (35 days = 840 hours)**. Note log axis. S2 error for derived S2 is approximately 9 cm for the VBayes at exact Envisat sampling period.

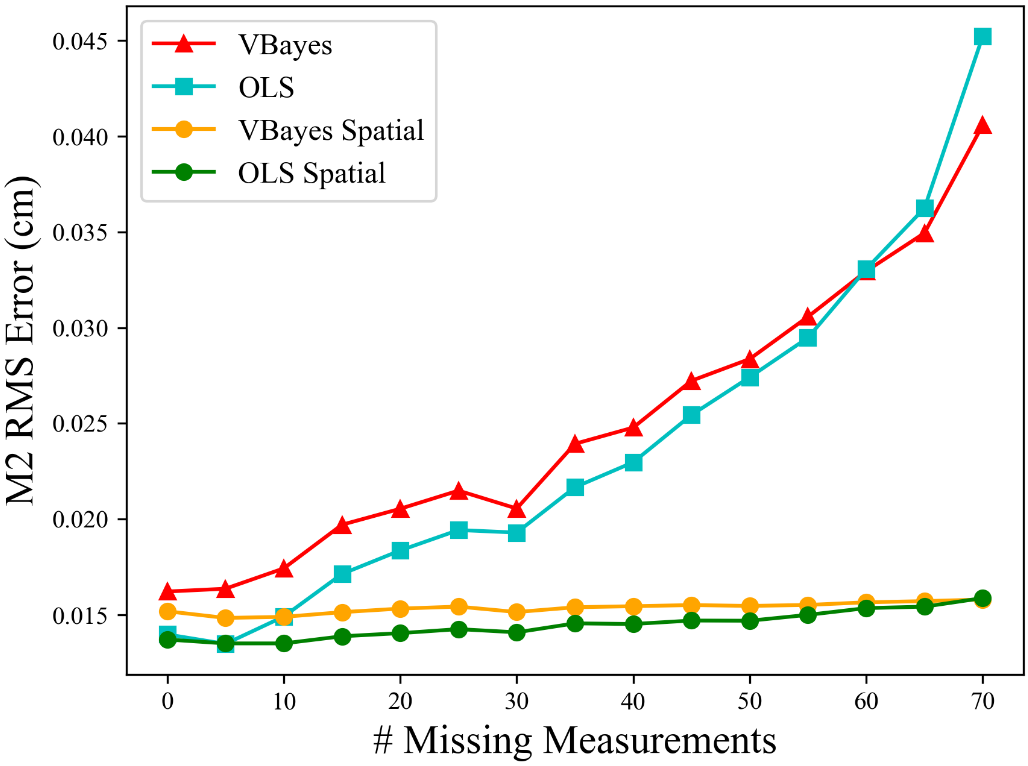

Figure A.3: **Comparisons of VBayes and OLS using spatial coherence.** Minimal noise case using only 100 $dB$ of WGN.

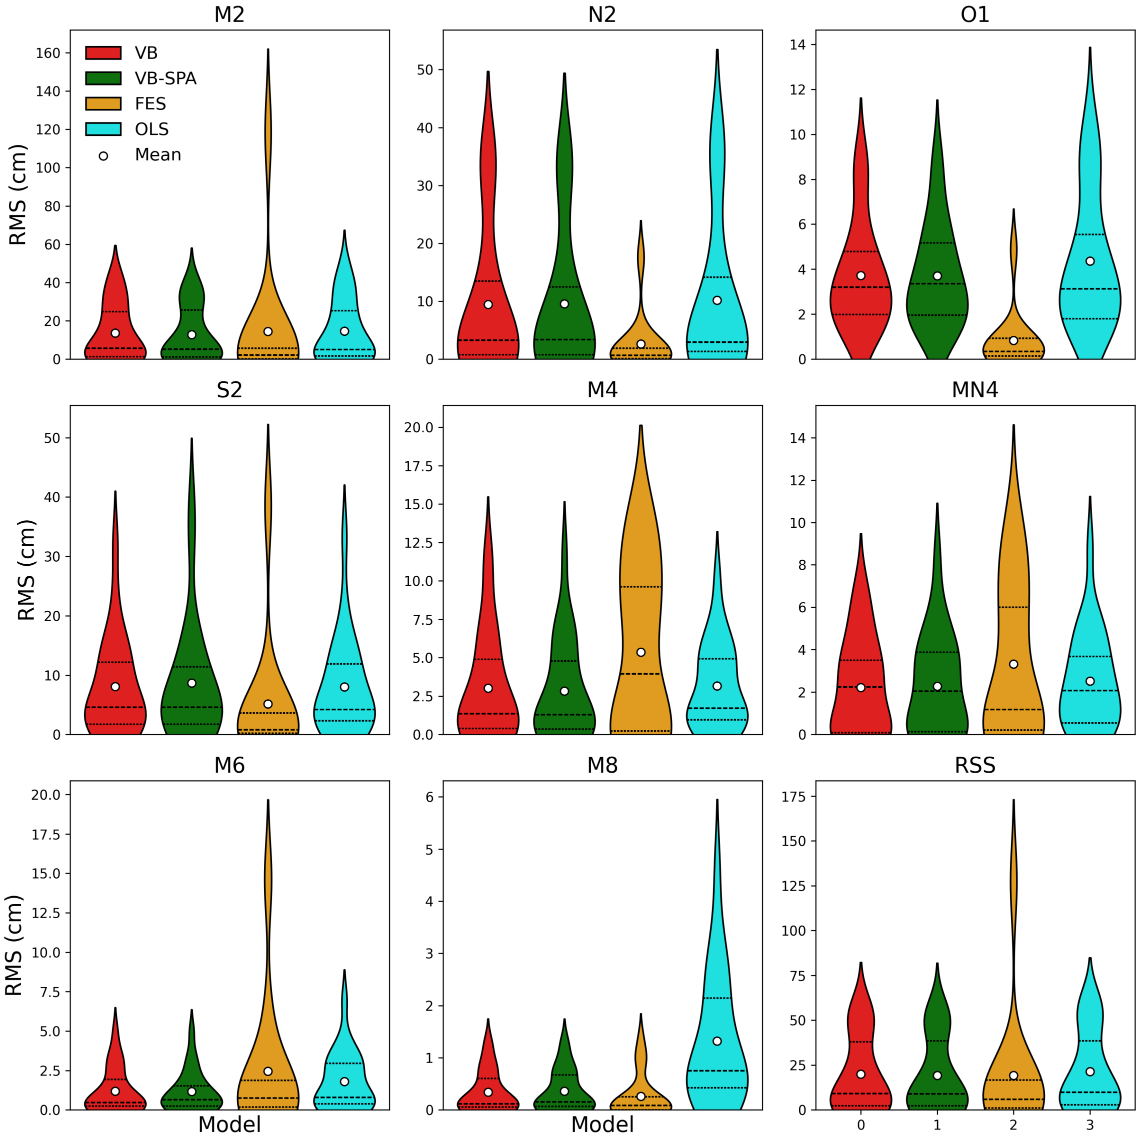

Figure A.4: **Comparisons of RMS model performance when performing direct empirical analysis of the sea-surface height.** (SLA + FES tide correction)

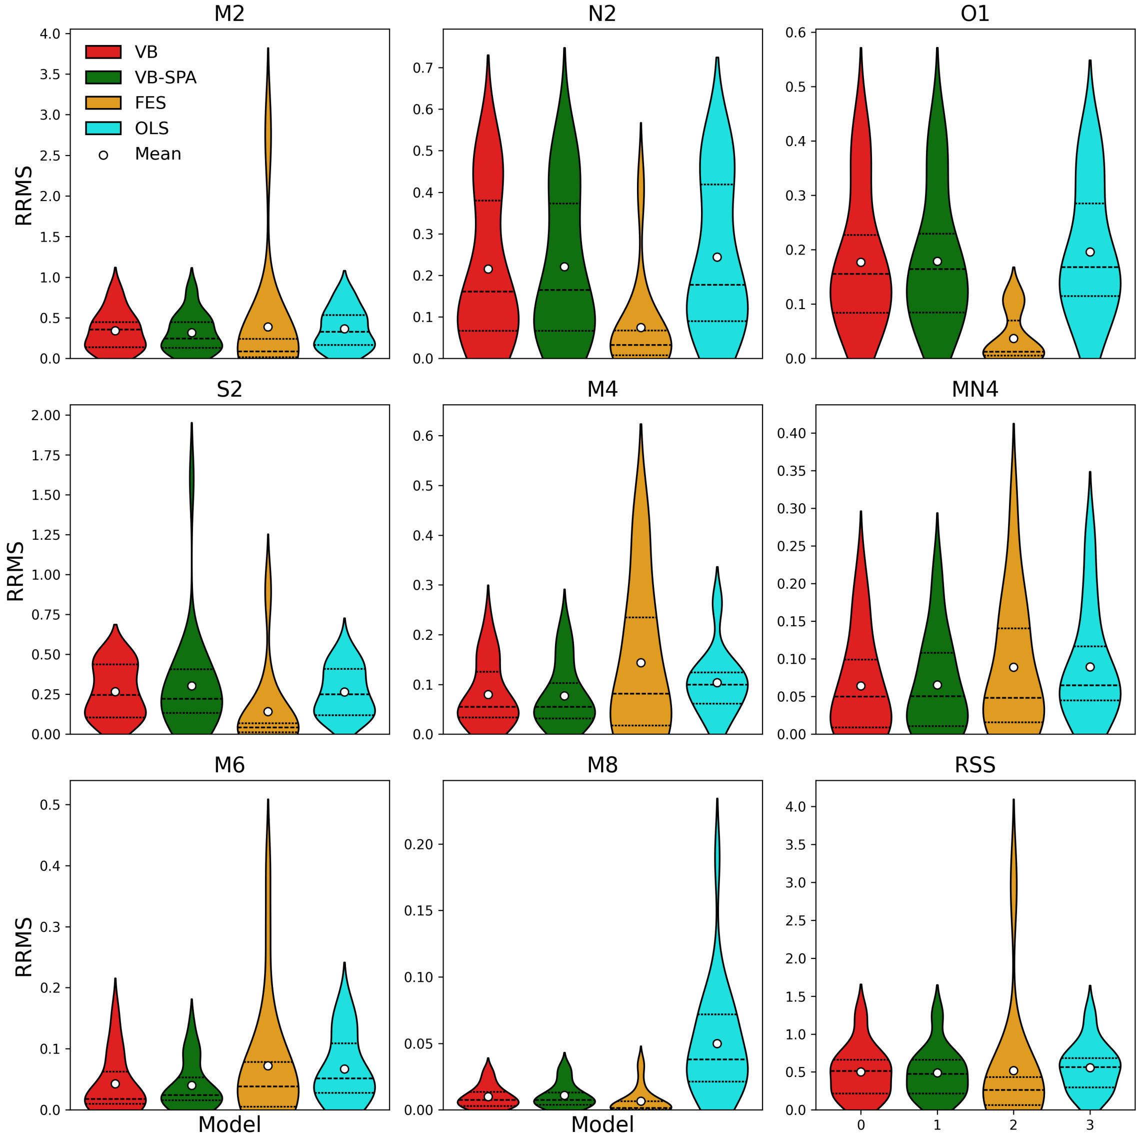

Figure A.5: **Comparisons of Relative RMS (RRMS) model performance when performing direct empirical analysis of the sea-surface height.** (SLA + FES tide correction)

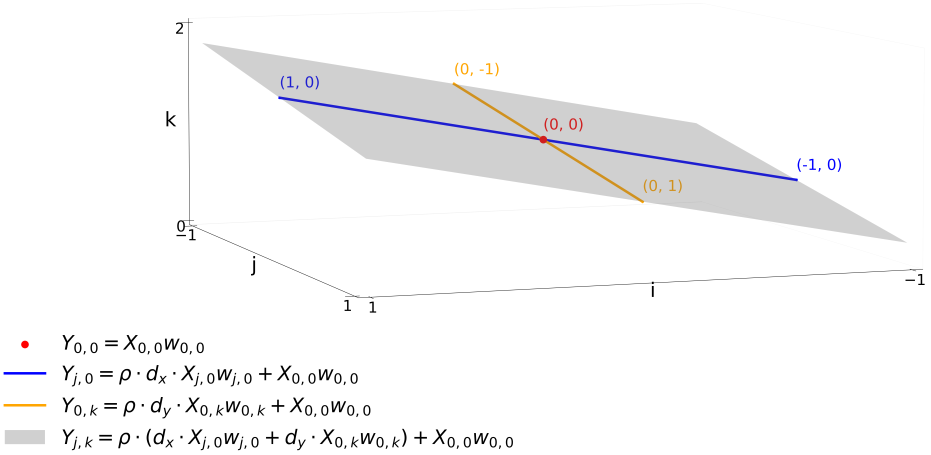

Figure A.6: **Spatially coherent inductive bias visualization** Relationship between different points.

# B Supporting information for Observations of the seiche that shook the world

### Contents of this appendix

1. Figure B1-B7
2. Tables B1-B2

Table B.1: Bounding box for SWOT querying

| Corner | Lat | Lon |
|---|---|---|
| Bottom Left | 72.77 | -27.51 |
| Top Left | 72.89 | -27.51 |
| Bottom Right | 72.77 | -26.42 |
| Top Right | 72.89 | -26.42 |

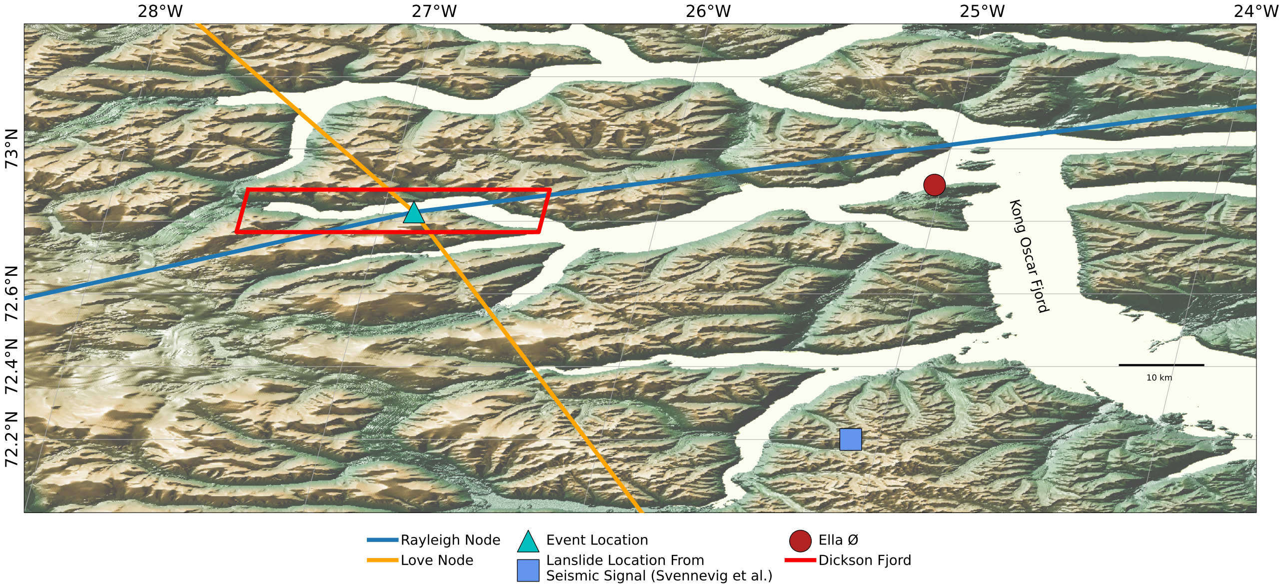

Figure B.1: Overview map of the study region including relevant adjacent locations. Background DEM is the Copernicus Global Digital Elevation Model (9).

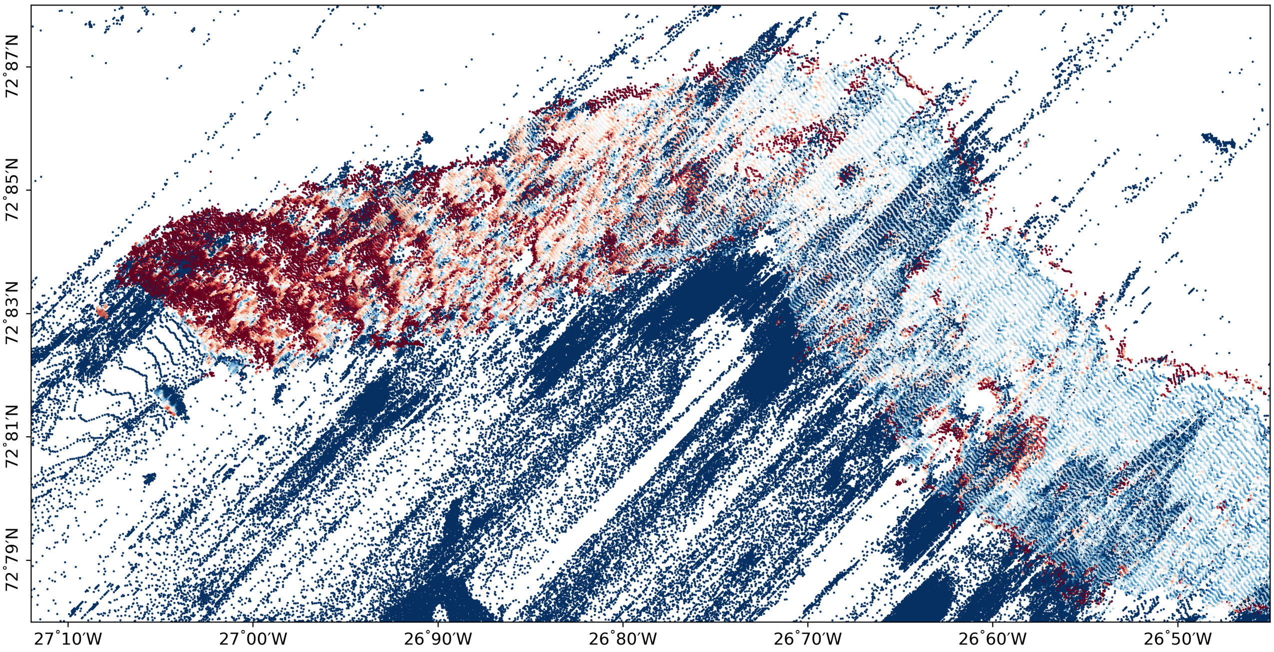

Figure B.2: Example SWOT observation with noise contamination. Data is processed using the same procedure described in Section 4.0.1. Observation produced by data file: SWOT_L2_HR_PIXC_005_292_021R_20231023T020832_20231023T020843_PGC0_01.

Table B.2: Tidal Constituents

| Constituent | Origin | Type | Period (hours) | CTD Amplitude (cm) |
|---|---|---|---|---|
| M2 | Principal lunar | semidiurnal | 12.42 | 43.19 |
| S2 | Principal solar | semidiurnal | 12.00 | 23.00 |
| N2 | Elliptical lunar | semidiurnal | 12.66 | 9.41 |
| K1 | Declinational luni-solar | diurnal | 23.93 | 8.06 |
| O1 | Principal lunar | diurnal | 25.82 | 7.49 |

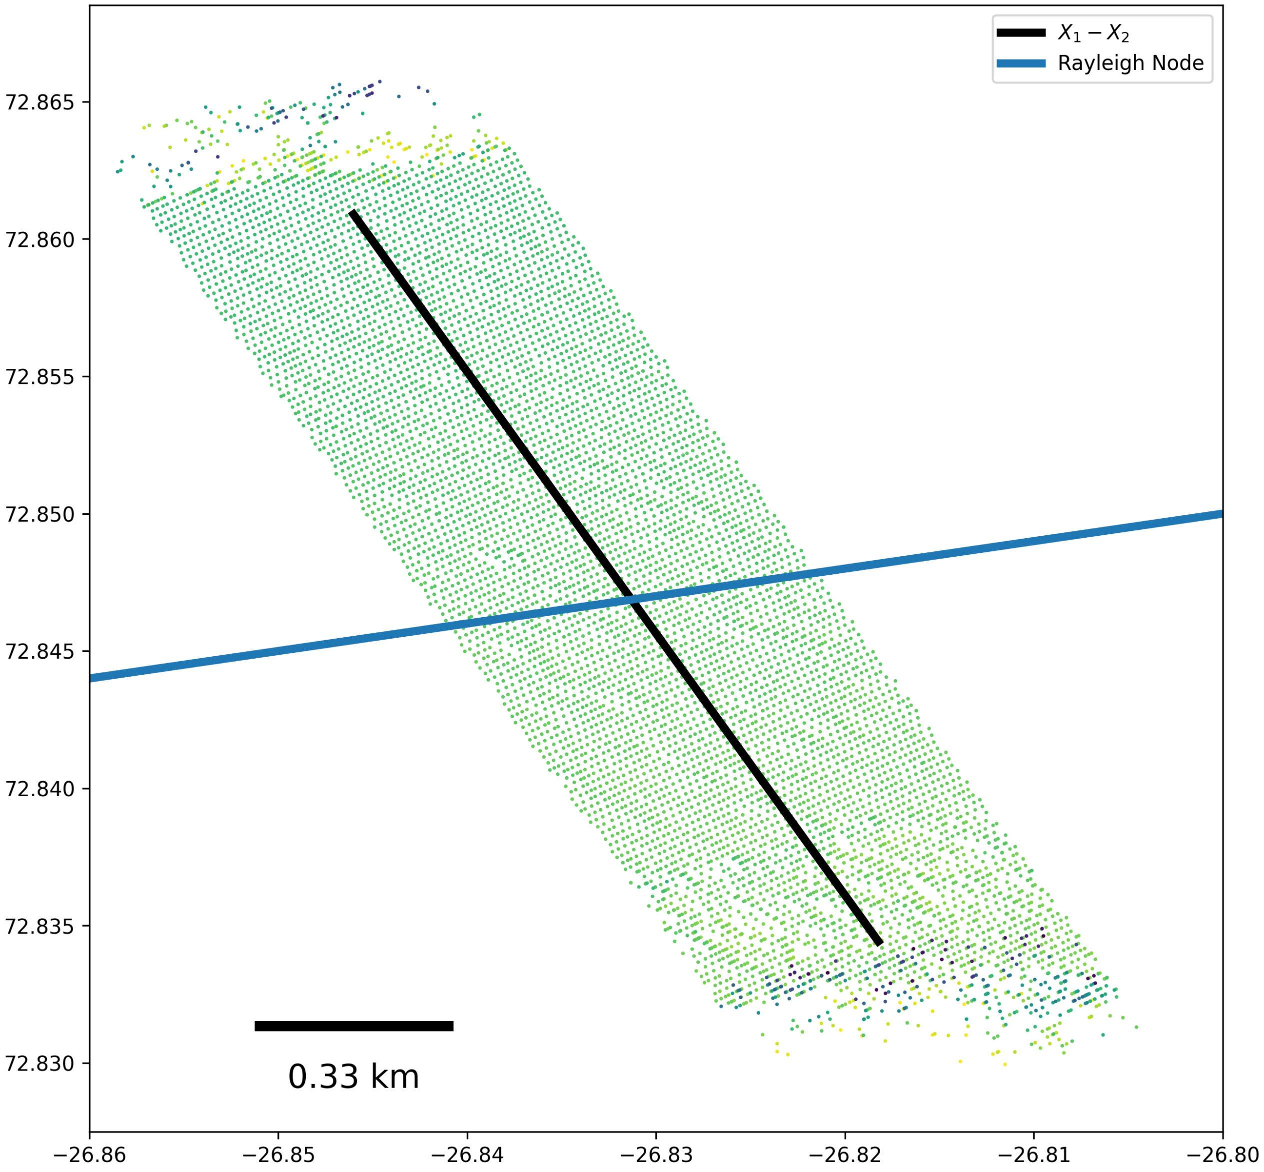

Figure B.3: Example SWOT cross-section used to compute the cross-channel slope. SWOT observations are colored according to sea-surface height.

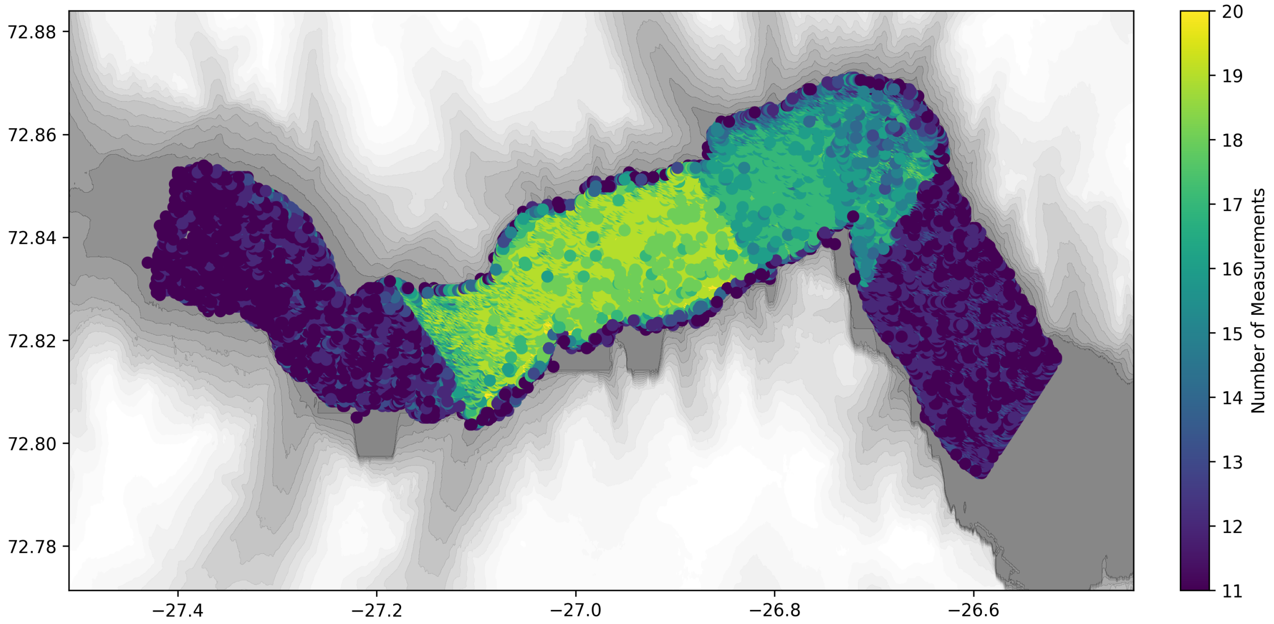

Figure B.4: Number of SWOT measurements available for tidal estimation.

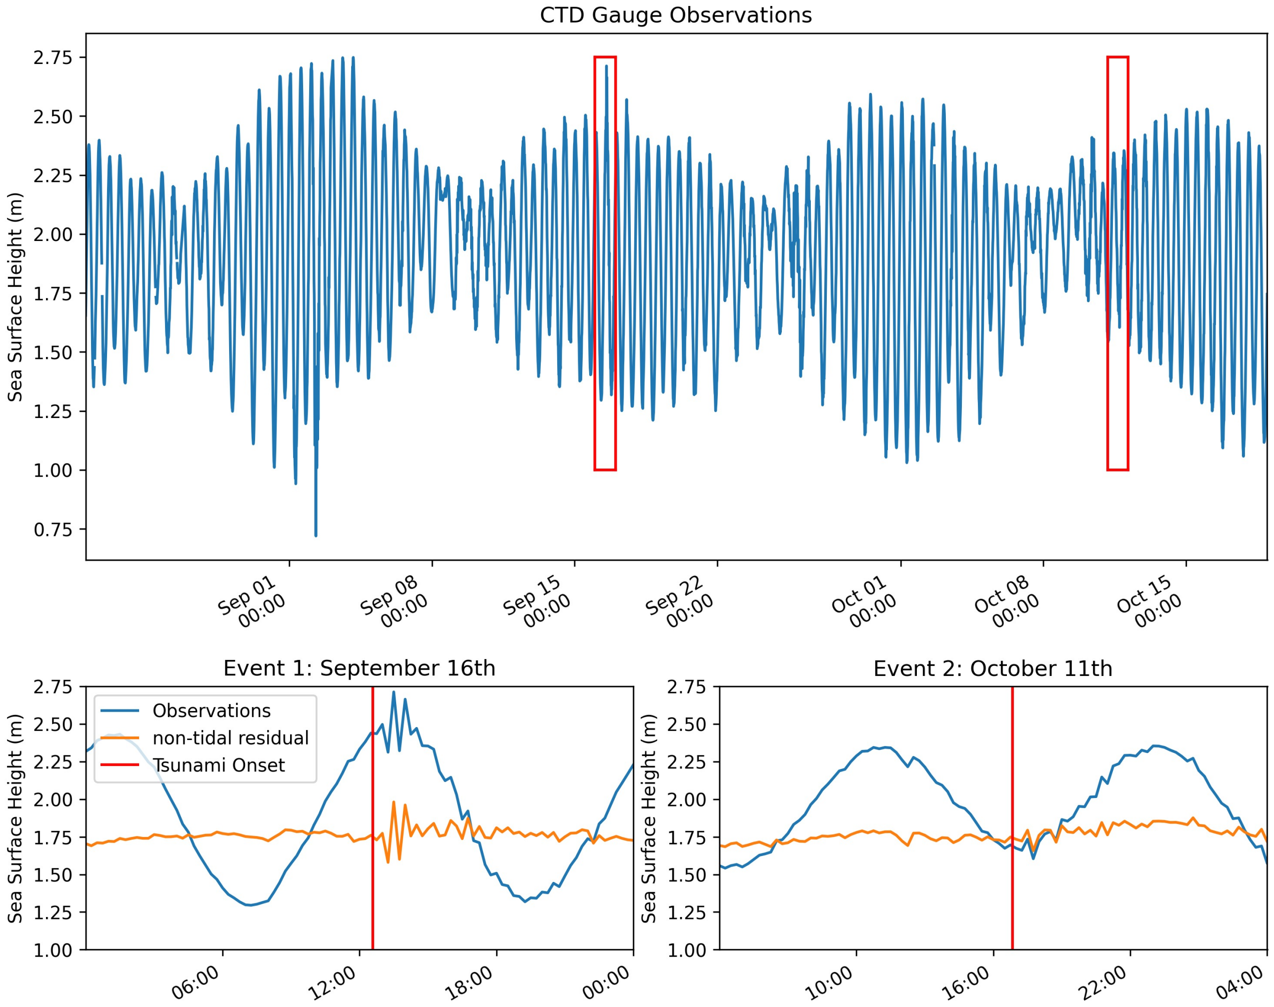

Figure B.5: Sea level observations from the CTD gauge located in the inner reaches of the Dickson fjord during the September and October events. Observations and non-tidal residuals are shown for the hours before and after tsunami onset. Observations are sampled at 15 minute intervals.

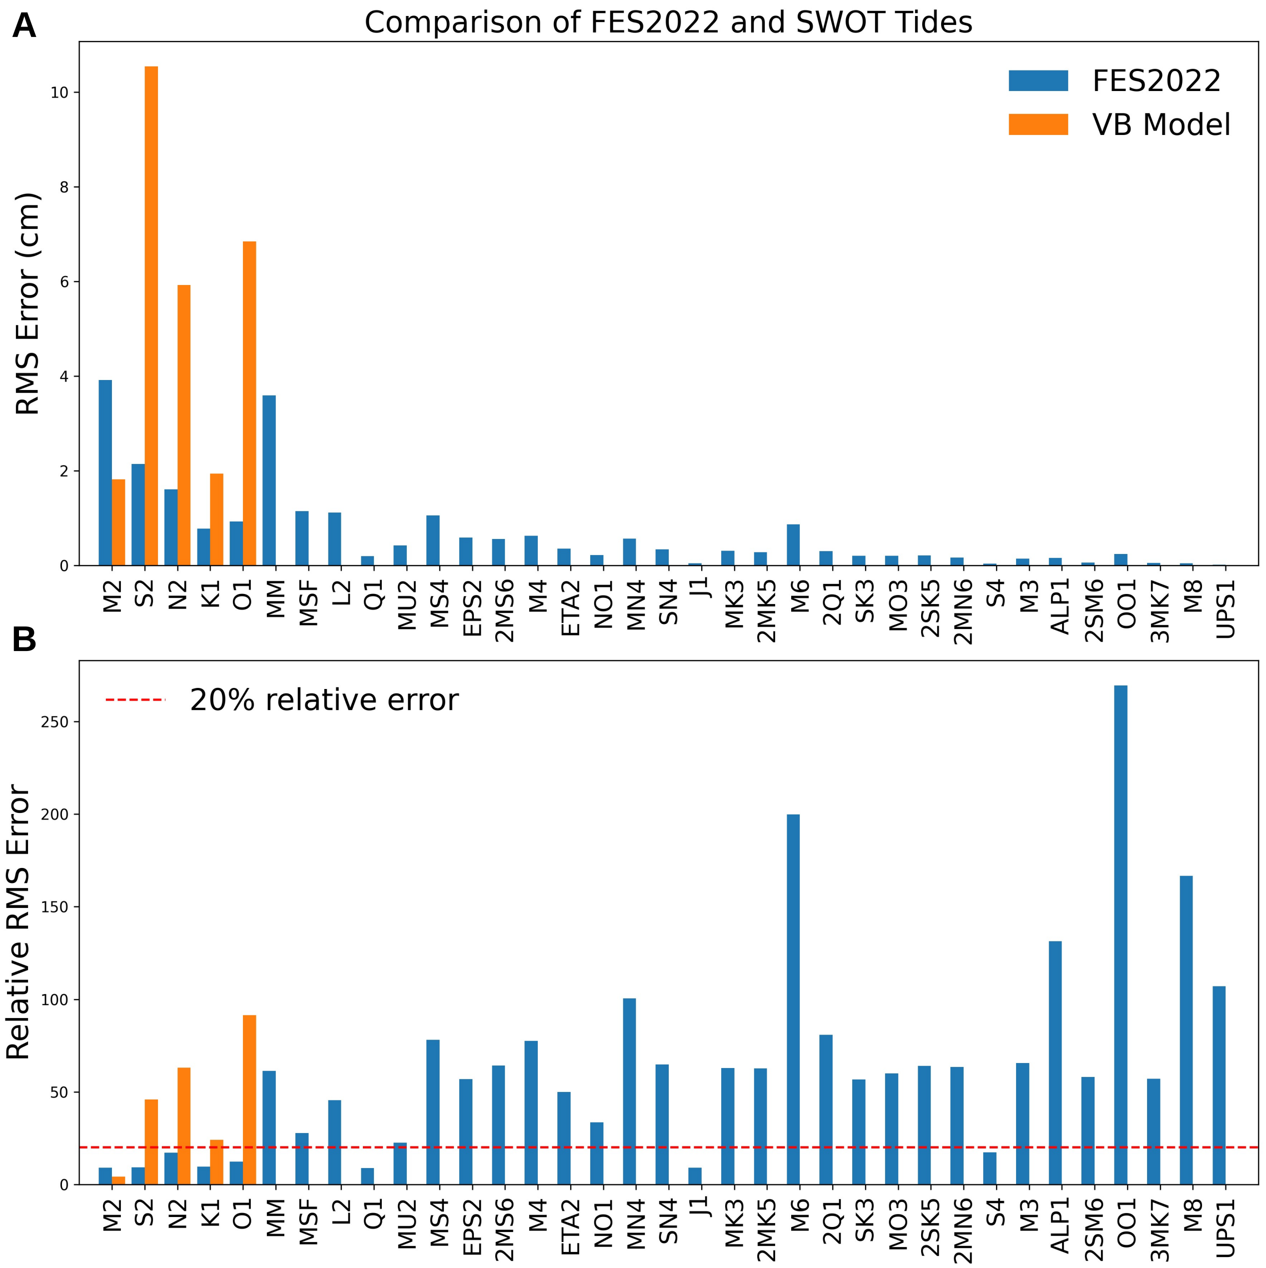

Figure B.6: Comparisons of FES2022 and VB derived tides from SWOT with the in-situ CTD Gauge. Upper Panel shows the RMS error defined in Section 4.4.2.5. Lower Panel shows the relative RMS error to the magnitude of the CTD gauge constituent. Note VB tides are only estimated for M2, N2, S2, K1, and O1.

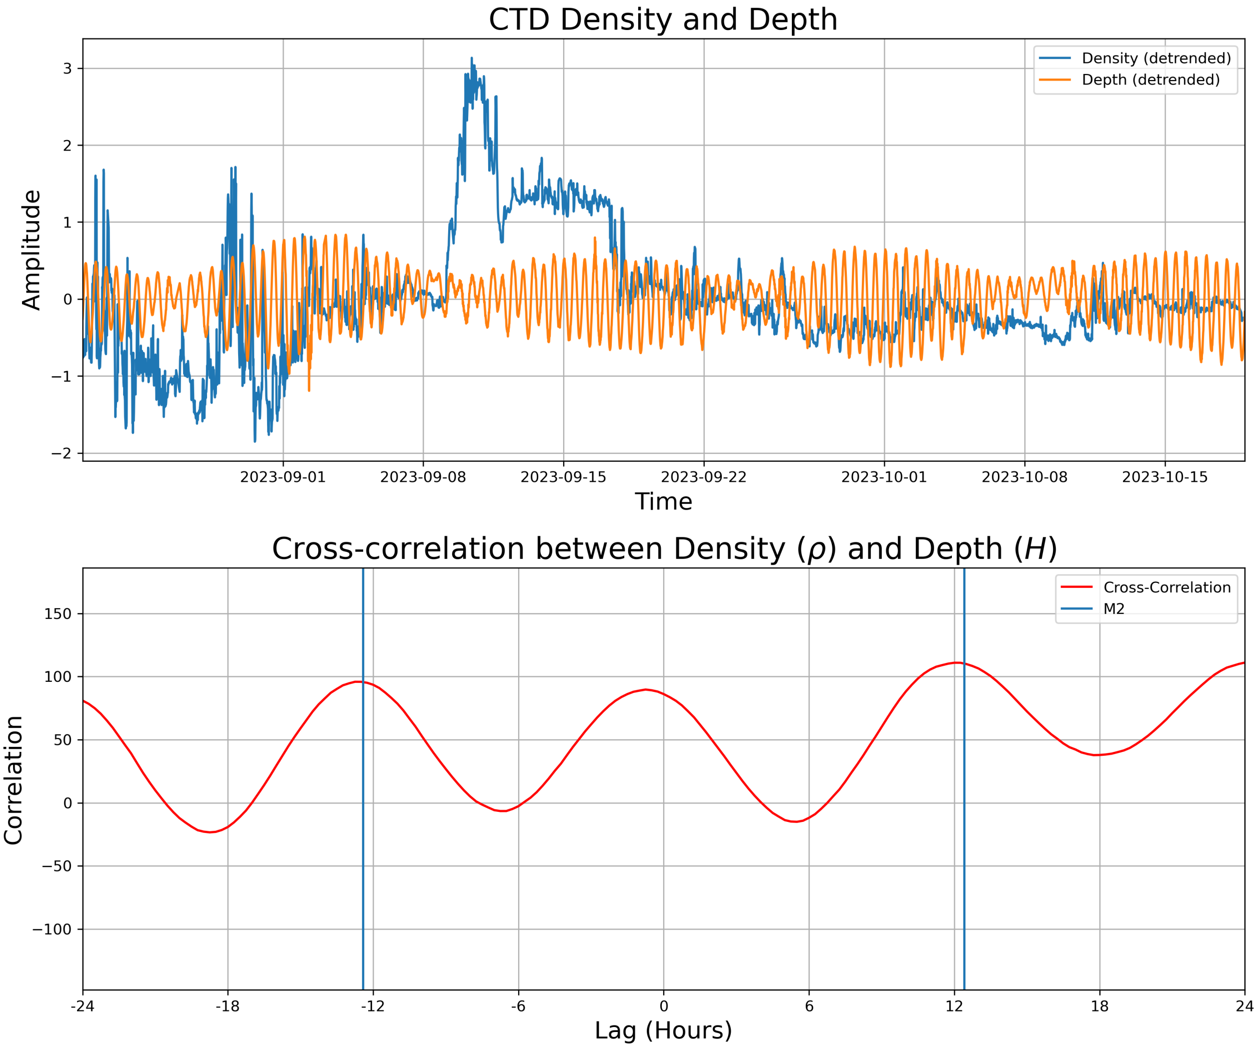

Figure B.7: Cross-Correlation between depth and density at Dickson fjord CTD Gauge. Top panel shows timeseries of detrended sea-surface height and density respectively. Lower panel shows cross-correlation.

# C Supporting information for RTide: Automating the tidal response method

## Contents of this appendix

1. Figures C1 to C4
2. Tables C1-C2

Figure C.3 shows the learned SHAP response distributions for the tidal river Dee in Chester, UK. The learned response distributions for the spherical harmonic Gravitational input functions are shown at various upstream river levels. Note that the superscript corresponds to the tidal species the input function drives (e.g. 1 = diurnal, 2 = semi-diurnal, etc.). The variance of the associated distributions is reported and clearly shows a damping of the sea-level response to the gravitational tidal forcing as the upstream river level increases.

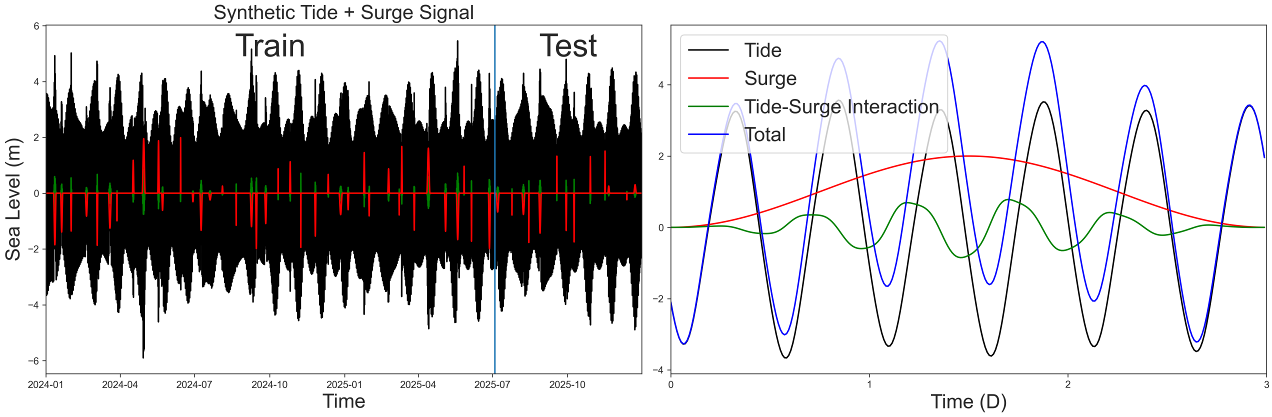

Figure C.1: **Visualization of the simulated tide, surge, and tide-surge signals used in Section 3.1.** Panel 1 shows the entire simulated signal comprised of 26 surge events. Panel 2 shows the evolution of a 3-day surge event.

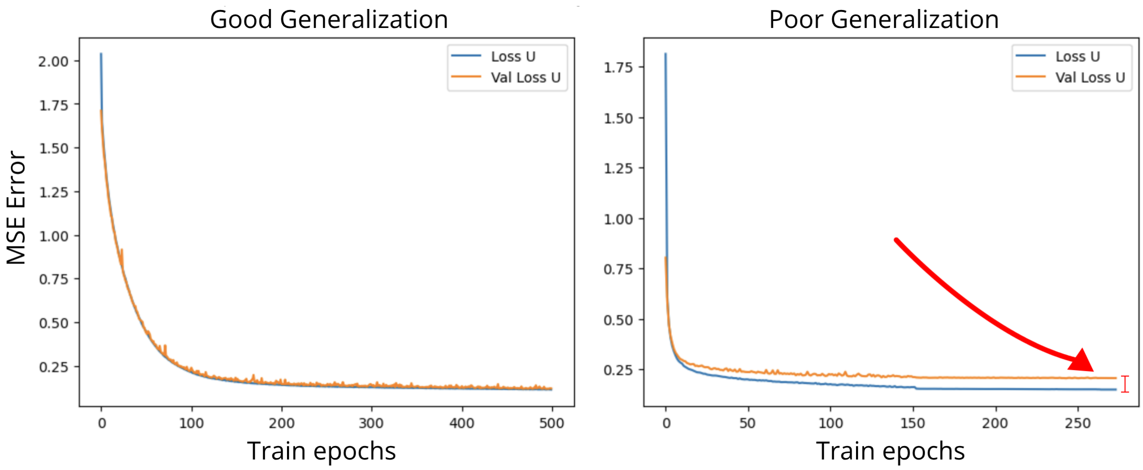

Figure C.2: **Comparisons of training loss curves for models with good generalization and bad generalization.** The separation between train set loss (Loss U), and the validation set (Val Loss U) indicates a poor performance out-of-sample and thus model overfitting. Typically this is usually a result of inappropriate input functions being used.

| Lags | Symmetrical | RMSE | MAE | Max Error | $R^2$ |
|---|---|---|---|---|---|
| $\tau$ = 2, s = 3 | False | 1.00988 | 0.76396 | 4.69893 | 0.67886 |
| $\tau$ = 2, s = 3 | True | 0.86403 | 0.65291 | 3.71645 | 0.76759 |
| Non-Uniform | False | 0.42070 | 0.33457 | 1.57420 | 0.95483 |
| Non-Uniform | True | 0.18183 | 0.13988 | 0.84289 | 0.99156 |

Table C.1: **Comparisons of RTide out of sample tide prediction accuracy based on selection of time-lags.** Non-Uniform refers to our proposed set of time-lags in Section 2.2.3. Note that this is the performance of these lags in our RTide implementation, which may differ from that of the conventional response method.

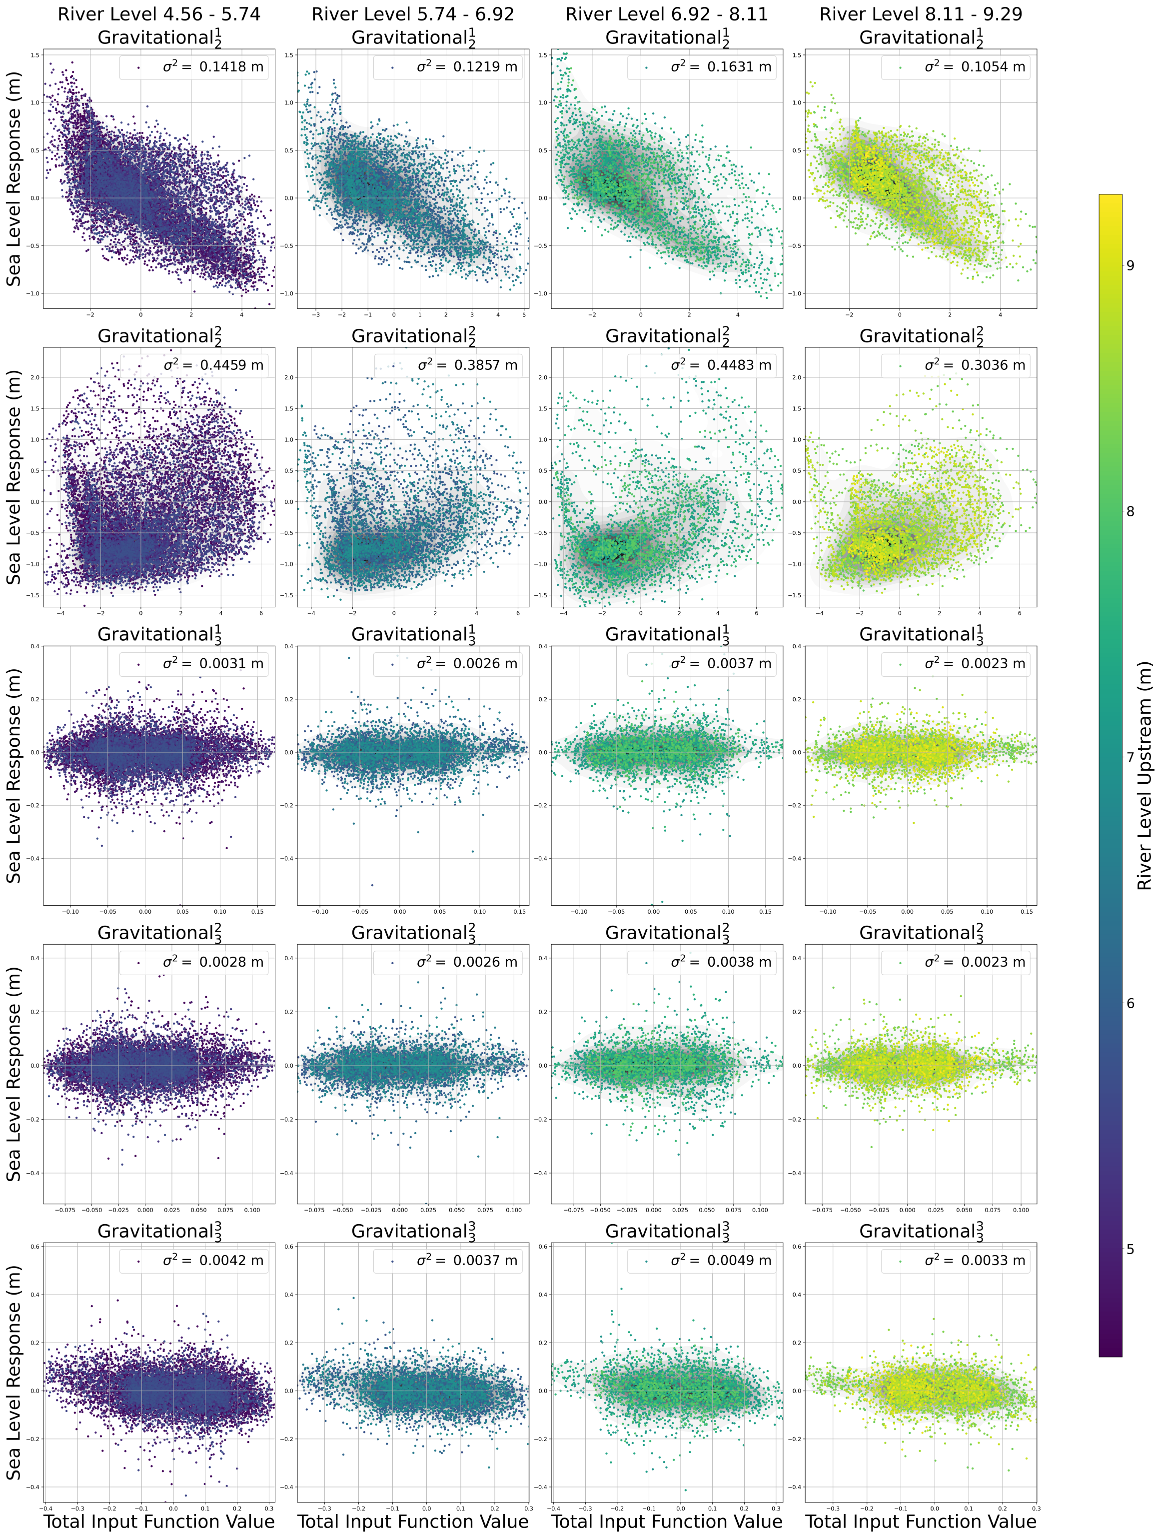

Figure C.3: **Comparisons of learned Response distributions for Gravitational input functions at Chester, UK based on upstream river-level**. Data taken from the year 2020. Here the subscript and superscript refer to the spherical harmonic degree and order respectively with the order separating the input functions according to species. Total input functions are obtained by summing the SHAP values of all time-lags of a given prediction. The variance of the distribution is reported and confirms the visual "squashing" seen in the response distributions at high river levels.

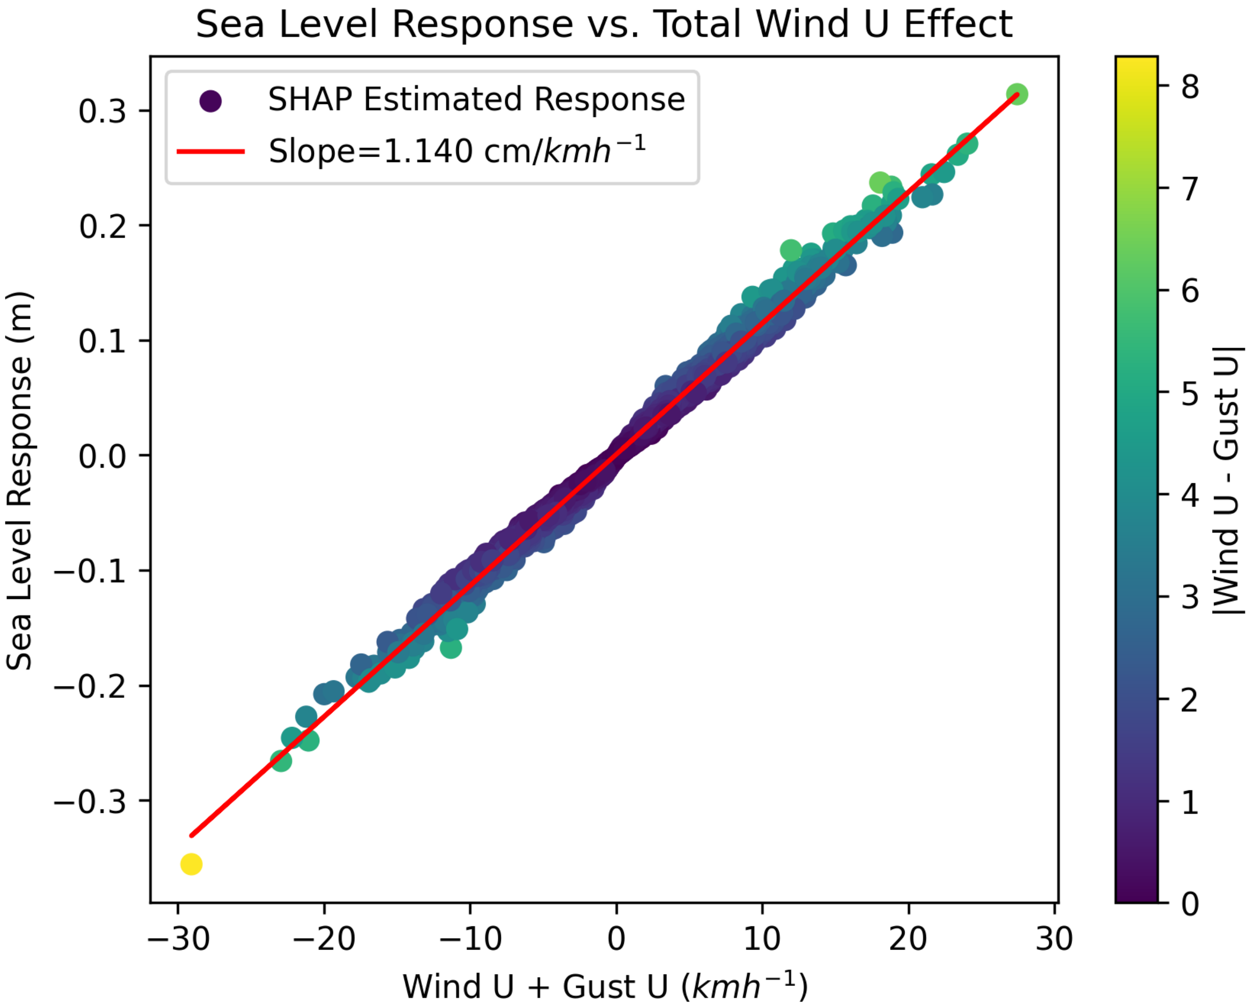

Figure C.4: **Learned SHAP dependence to realtime wind forcing in U direction by summing sustained and gust components from Money Point, VA, USA.** Responses are estimated for first 6 months of 2016 after being trained on data from the year 2014. Data points are colored based on the absolute difference between the sustained and gust components.

| Constituent | RTide | OLS HA | IRLS HA |
|---|---|---|---|
| M2 | 0.695704 | 0.700515 | 0.692926 |
| N2 | 0.732647 | 0.742326 | 0.741203 |
| S2 | 0.341145 | 0.344736 | 0.342417 |
| L2 | 0.113816 | 0.127491 | 0.128792 |
| K1 | 0.145476 | 0.145122 | 0.144331 |
| NU2 | 0.141982 | 0.148616 | 0.148055 |
| O1 | 0.128958 | 0.136669 | 0.100519 |
| K2 | 0.104260 | 0.105884 | 0.124208 |
| 2N2 | 0.073575 | 0.074250 | 0.075066 |
| M4 | 0.054395 | 0.078930 | 0.054844 |
| LDA2 | 0.056129 | 0.066950 | 0.052076 |
| P1 | 0.064692 | 0.069703 | 0.066135 |
| M6 | 0.058307 | 0.051206 | 0.060394 |
| 2MN6 | 0.017841 | 0.041493 | 0.030160 |
| SSA | 0.011102 | 0.047023 | 0.034311 |
| MM | 0.031390 | 0.010564 | 0.021471 |
| MN4 | 0.018754 | 0.027909 | 0.031617 |
| 2MS6 | 0.029002 | 0.027133 | 0.027643 |
| Q1 | 0.025232 | 0.003016 | 0.021439 |
| MS4 | 0.009766 | 0.008117 | 0.007360 |
| OQ2 | 0.001907 | 0.006553 | 0.011235 |
| MF | 0.021641 | 0.018951 | 0.006324 |
| EPS2 | 0.017846 | 0.022367 | 0.012020 |
| NO1 | 0.010148 | 0.018463 | 0.013559 |
| MU2 | 0.016172 | 0.011659 | 0.013863 |
| MSF | 0.009547 | 0.013468 | 0.009951 |
| MSN2 | 0.005683 | 0.005402 | 0.011039 |
| MKS2 | 0.009707 | 0.009167 | 0.010309 |
| MSM | 0.012011 | 0.008740 | 0.011972 |
| TAU1 | 0.001014 | 0.002285 | 0.003858 |
| CHI1 | 0.000923 | 0.007188 | 0.000115 |
| 2MK6 | 0.008420 | 0.006492 | 0.002106 |
| RSS | 1.115014 | 1.131959 | 1.121273 |

Table C.2: Comparison of RMS error (cm) in the derived harmonic constituent between RTide, Ordinary Least Squares (OLS) HA, and Iteratively Reweighted Least Squares (IRLS) HA. Only the 32 largest constituents are shown. The Root Sum of Squares (RSS) is computed for the set of all 59 constituents used to generate the synthetic tidal signal. RTide models were trained with atmospheric pressure input as an additional input function.

# D Supporting information for Response-Based prediction of tidal currents

## Contents of this appendix

1. Figures D1
2. Table D1-D2

| Node | Number | Latitude | Longitude |
|---|---|---|---|
| A | 500 | $59.8599068915N^{\circ}$ | $4.9797111830W^{\circ}$ |
| B | 4768 | $58.6557103353N^{\circ}$ | $3.1391667722W^{\circ}$ |
| C | 4496 | $58.6918591258N^{\circ}$ | $3.1296765346W^{\circ}$ |

Table D.1: Location of Case Study Nodes at Pentland Firth.

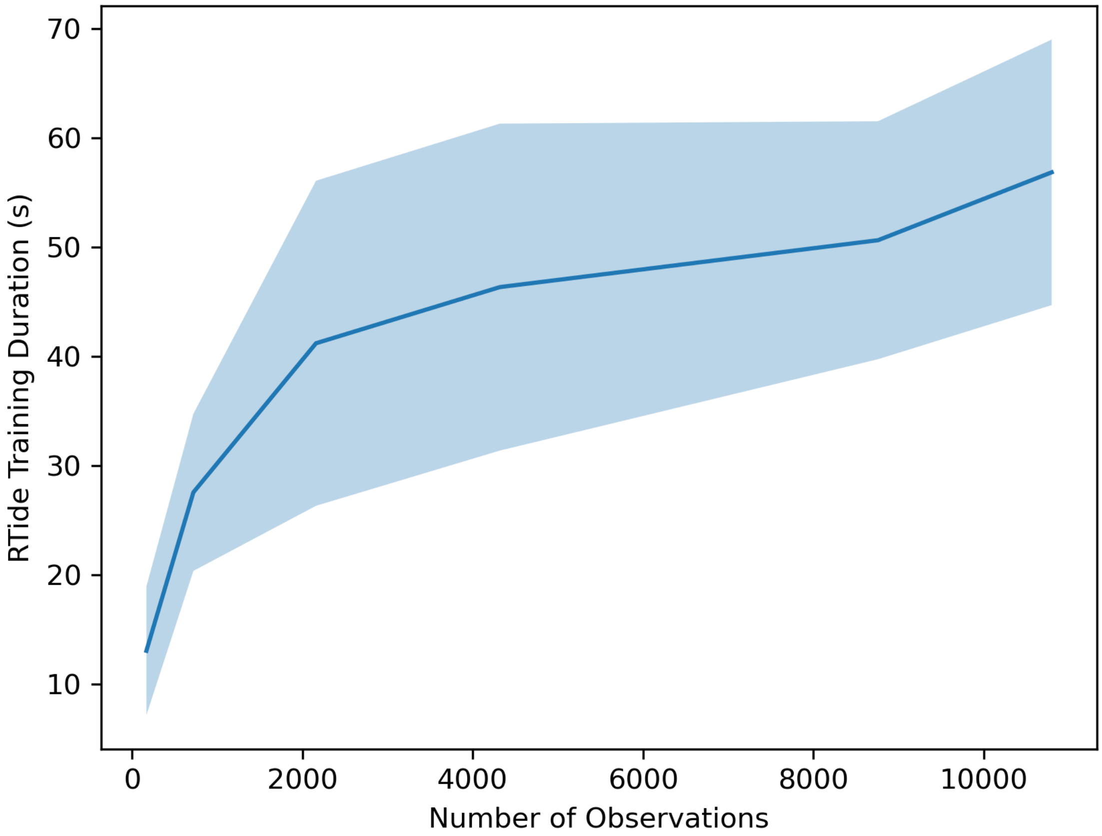

Figure D.1: **RTide training time on single Apple M3 CPU as a function of the number of observations used for training. Standard deviations are computed across 10 runs. Performance will vary by hardware, but significant speedups can be achieved if GPUs are available.**

| Constituent | RMS_rtide | RMS_utide | Actual |
|---|---|---|---|
| M2 | 0.002736 | 0.005684 | 1.548235 |
| S2 | 0.036669 | 0.003219 | 0.527180 |
| N2 | 0.006481 | 0.017489 | 0.297678 |
| K2 | 0.014814 | NaN | 0.115721 |
| MU2 | 0.002259 | 0.007527 | 0.093891 |
| M4 | 0.003023 | 0.003252 | 0.077345 |
| 2MS6 | 0.003842 | 0.003246 | 0.071057 |
| NU2 | 0.006673 | NaN | 0.065973 |
| M6 | 0.007071 | 0.008095 | 0.066981 |
| MSF | 0.010238 | 0.002148 | 0.064909 |
| L2 | 0.004945 | 0.009936 | 0.062232 |
| O1 | 0.005652 | 0.005126 | 0.052799 |
| K1 | 0.000944 | 0.004332 | 0.051445 |
| MS4 | 0.002382 | 0.003230 | 0.049760 |
| 3MS8 | 0.001230 | NaN | 0.045876 |
| 2MN6 | 0.003580 | 0.003042 | 0.042798 |
| LDA2 | 0.005439 | NaN | 0.041724 |
| EPS2 | 0.010243 | 0.005512 | 0.036572 |
| MM | 0.001020 | 0.001773 | 0.031499 |
| M8 | 0.004727 | 0.002858 | 0.029452 |
| MN4 | 0.001368 | 0.001806 | 0.029608 |
| MSN2 | 0.010581 | NaN | 0.023877 |
| 2N2 | 0.014536 | NaN | 0.023637 |
| 2SM6 | 0.002424 | 0.001961 | 0.022017 |
| MKS2 | 0.011095 | NaN | 0.017797 |
| 3MN8 | 0.003988 | NaN | 0.016219 |
| 2MK6 | 0.002905 | NaN | 0.015379 |
| MF | 0.003157 | NaN | 0.013259 |

| | | | |
|---|---|---|---|
| MK3 | 0.000813 | 0.000519 | 0.012564 |
| MSK6 | 0.001111 | NaN | 0.011913 |
| MK4 | 0.001506 | NaN | 0.012564 |
| TAU1 | 0.003340 | NaN | 0.011164 |
| GAM2 | 0.020118 | NaN | 0.010410 |
| M12 | 0.004779 | NaN | 0.010038 |
| SIG1 | 0.003852 | NaN | 0.009274 |
| 2MK5 | 0.002671 | 0.002968 | 0.009352 |
| P1 | 0.005327 | NaN | 0.008944 |
| MO3 | 0.003570 | 0.004661 | 0.009239 |
| 3MK7 | 0.002239 | 0.003126 | 0.008778 |
| SO1 | 0.003283 | NaN | 0.007891 |
| OQ2 | 0.006577 | NaN | 0.007700 |
| SO3 | 0.004474 | NaN | 0.007184 |
| SK4 | 0.002868 | NaN | 0.007189 |
| M10 | 0.003617 | NaN | 0.006467 |
| SK3 | 0.006415 | 0.005045 | 0.005852 |
| H1 | 0.004614 | NaN | 0.005559 |
| S4 | 0.002350 | 0.002103 | 0.005255 |
| SSA | 0.007116 | NaN | 0.004687 |
| NO1 | 0.002456 | 0.000797 | 0.004310 |
| MSM | 0.001583 | NaN | 0.004341 |
| RHO1 | 0.003020 | NaN | 0.004268 |
| SN4 | 0.005434 | 0.001782 | 0.004596 |
| Q1 | 0.001632 | 0.002397 | 0.004248 |
| THE1 | 0.001467 | NaN | 0.003880 |
| ETA2 | 0.002014 | 0.004704 | 0.003562 |
| N4 | 0.003391 | NaN | 0.003450 |
| M3 | 0.005094 | 0.003962 | 0.003218 |
| OO1 | 0.000353 | 0.003819 | 0.003188 |

| | | | |
|---|---|---|---|
| H2 | 0.005117 | NaN | 0.002750 |
| J1 | 0.006684 | 0.000727 | 0.002793 |
| S1 | 0.001146 | NaN | 0.002713 |
| R2 | 0.004257 | NaN | 0.002584 |
| PI1 | 0.001721 | NaN | 0.002329 |
| T2 | 0.041185 | NaN | 0.002287 |
| PSI1 | 0.001437 | NaN | 0.001851 |
| BET1 | 0.001050 | NaN | 0.001933 |
| PHI1 | 0.001432 | NaN | 0.001781 |
| CHI1 | 0.005539 | NaN | 0.001914 |
| ALP1 | 0.002441 | 0.004175 | 0.001820 |
| 2Q1 | 0.003201 | 0.004277 | 0.001595 |
| SA | 0.004159 | NaN | 0.001575 |
| 2SK5 | 0.000888 | 0.003795 | 0.001438 |
| UPS1 | 0.001987 | 0.002720 | 0.000507 |

Table D.2: Constituent RMS values computed from RTide and UTide models using 180 days of synthetic Pentland Firth training data. Equivalent to Figure 3. Note for this table the RTide constituents are directly estimated using harmonic analysis to illustrate if the model is learning spurious constituents. NaN values are given for the UTide estimates if constituents are not derivable beyond a Rayleigh criterion of 1.

# References

[1] W. Munk, K. Hasselmann, Super-resolution of tides, Studies on oceanography (1964) 339–344.

[2] Met Office, Cartopy: a cartographic python library with a Matplotlib interface, Exeter, Devon (2010 - 2015).
URL `http://scitools.org.uk/cartopy`

[3] S. I. of Oceanography, Global seismograph network-iris/ida, International Federation of Digital Seismograph Networks.

[4] A. S. L. (ASL)/USGS, Global seismograph network (gsn-iris/usgs), International Federation of Digital Seismograph Networks.

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

[6] S. Lundberg, A unified approach to interpreting model predictions, arXiv preprint arXiv:1705.07874.

[7] M. R. Klema, A. G. Pirzado, S. K. Venayagamoorthy, T. K. Gates, Analysis of acoustic doppler current profiler mean velocity measurements in shallow flows, Flow Measurement and Instrumentation 74 (2020) 101755.

[8] M. Liu, W. Li, C. Wang, R. Billinton, J. Yu, Reliability evaluation of a tidal power generation system considering tidal current speeds, IEEE Transactions on Power Systems 31 (4) (2015) 3179–3188.

[9] European Space Agency, Copernicus global digital elevation model, distributed by OpenTopography. Accessed: 2024-10-10 (2024).
URL `https://doi.org/10.5069/G9028PQB`

[10] B. B. Parker, Tidal analysis and prediction.

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

[12] G. Godin, The Resolution of Tidal Constituents, The International Hydrographic Review.

[13] T. Monahan, T. Tang, S. Roberts, T. A. Adcock, Tidal and mean sea surface corrections from and for swot using a spatially coherent variational bayesian harmonic analysis, Authorea Preprints.

[14] T. Monahan, T. Tang, S. Roberts, T. A. Adcock, First observations of the seiche that shook the world, arXiv preprint arXiv:2411.02469.

[15] P. Matte, D. A. Jay, E. D. Zaron, Adaptation of classical tidal harmonic analysis to nonstationary tides, with application to river tides, Journal of Atmospheric and Oceanic Technology 30 (3) (2013) 569–589.

[16] D. E. Cartwright, 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) (1968) 1–55.

[17] T. Monahan, T. Tang, S. Roberts, T. Adcock, Response framework: Tidal analysis and prediction through physics-informed ml.

[18] R. A. Flather, Existing operational oceanography, Coastal Engineering 41 (1-3) (2000) 13–40.

[19] M. G. Hart-Davis, G. Piccioni, D. Dettmering, C. Schwatke, M. Passaro, F. Seitz, Eot20: A global ocean tide model from multi-mission satellite altimetry, Earth System Science Data 13 (8) (2021) 3869–3884.

[20] T. Monahan, T. Tang, T. A. Adcock, A hybrid model for online short-term tidal energy forecasting, Applied Ocean Research 137 (2023) 103596.

[21] W. Munk, Solving the tidal problem (2000).

[22] F. H. Lyard, D. J. Allain, M. Cancet, L. Carrère, N. Picot, Fes2014 global ocean tide atlas: design and performance, Ocean Science 17 (3) (2021) 615–649.

[23] D. Stammer, R. Ray, O. B. Andersen, B. Arbic, W. Bosch, L. Carrère, Y. Cheng, D. Chinn, B. Dushaw, G. Egbert, et al., Accuracy assessment of global barotropic ocean tide models, Reviews of Geophysics 52 (3) (2014) 243–282.

[24] D. A. Greenberg, F. Dupont, F. H. Lyard, D. R. Lynch, F. E. Werner, Resolution issues in numerical models of oceanic and coastal circulation, Continental Shelf Research 27 (9) (2007) 1317–1343.

[25] G. D. Egbert, R. D. Ray, B. G. Bills, Numerical modeling of the global semidiurnal tide in the present day and in the last glacial maximum, Journal of Geophysical Research: Oceans 109 (C3).

[26] G. D. Egbert, R. D. Ray, Significant dissipation of tidal energy in the deep ocean inferred from satellite altimeter data, Nature 405 (6788) (2000) 775–778.

[27] M. Tonani, P. Sykes, R. R. King, N. McConnell, A.-C. Péquignet, E. O'Dea, J. A. Graham, J. Polton, J. Siddorn, The impact of a new high-resolution ocean model on the met office north-west european shelf forecasting system, Ocean Science 15 (4) (2019) 1133–1158.

[28] M. Hart-Davis, O. Andersen, R. Ray, E. Zaron, C. Schwatke, R. Arildsen, D. Dettmering, Tides in complex coastal regions: early case studies from wide-swath swot measurements, Authorea Preprints.

[29] L. Carrere, F. Lyard, M. Cancet, D. Allain, M.-L. Dabat, E. Fouchet, E. Sahuc, Y. Faugere, G. Dibarboure, N. Picot, A new barotropic tide model for global ocean: Fes2022, in: 2022 Ocean Surface Topography Science Team Meeting, 2022, p. 43.

[30] A. F. Bennett, Inverse methods in physical oceanography, Cambridge university press, 1992.

[31] P. Schureman, Manual of harmonic analysis and prediction of tides, Vol. 4, US Government Printing Office, 1958.

[32] R. Pawlowicz, B. Beardsley, S. Lentz, Classical tidal harmonic analysis including error estimates in matlab using t_tide, Computers & geosciences 28 (8) (2002) 929–937.

[33] K. E. Leffler, D. A. Jay, Enhancing tidal harmonic analysis: Robust (hybrid l1/l2) solutions, Continental Shelf Research 29 (1) (2009) 78–88.

[34] D. L. Codiga, Unified tidal analysis and prediction using the UTide Matlab functions.

[35] K. George, B. Simon, The species concordance method of tide prediction in estuaries, The International Hydrographic Review.

[36] J. Williams, M. Irazoqui Apecechea, A. Saulter, K. J. Horsburgh, Radiational tides: their double-counting in storm surge forecasts and contribution to the highest astronomical tide, Ocean Science 14 (5) (2018) 1057–1068.

[37] H. Landau, Sampling, data transmission, and the nyquist rate, Proceedings of the IEEE 55 (10) (1967) 1701–1706.

[38] M. Mishali, Y. C. Eldar, From theory to practice: Sub-nyquist sampling of sparse wideband analog signals, IEEE Journal of selected topics in signal processing 4 (2) (2010) 375–391.

[39] D. E. Cartwright, R. Ray, Oceanic tides from geosat altimetry, Journal of Geophysical Research: Oceans 95 (C3) (1990) 3069–3090.

[40] O. B. Andersen, S. K. Rose, M. G. Hart-Davis, Polar ocean tides—revisited using cryosat-2, Remote Sensing 15 (18) (2023) 4479.

[41] B. Zetler, D. Cartwright, S. Berkman, Some comparisons of response and harmonic tide predictions, The International Hydrographic Review.

[42] W. Munk, B. Zetler, G. Groves, Tidal cusps, Geophysical Journal International 10 (2) (1965) 211–219.

[43] M. Amin, G. Lennon, On anomalous tides in australian waters, The International Hydrographic Review.

[44] E. Flinchem, D. Jay, An introduction to wavelet transform tidal analysis methods, Estuarine, Coastal and Shelf Science 51 (2) (2000) 177–200.

[45] M. Lobo, D. A. Jay, S. Innocenti, S. A. Talke, S. L. Dykstra, P. Matte, Implementing super-resolution of nonstationary tides with wavelets: An introduction to cwt_multi, Journal of Atmospheric and Oceanic Technology 41 (10) (2024) 969–989.

[46] S. Abdalla, A. A. Kolahchi, M. Ablain, S. Adusumilli, S. A. Bhowmick, E. Alou-Font, L. Amarouche, O. B. Andersen, H. Antich, L. Aouf, et al., Altimetry for the future: Building on 25 years of progress, Advances in Space Research 68 (2) (2021) 319–363.

[47] G. D. Egbert, R. D. Ray, Tidal prediction.

[48] R. Morrow, L.-L. Fu, F. Ardhuin, M. Benkiran, B. Chapron, E. Cosme, F. d'Ovidio, J. T. Farrar, S. T. Gille, G. Lapeyre, et al., Global observations of fine-scale ocean surface topography with the surface water and ocean topography (SWOT) mission, Frontiers in Marine Science 6 (2019) 232.

[49] P. L. Woodworth, J. M. Vassie, Reanalyses of maskelyne's tidal data at st. helena in 1761, Earth System Science Data 14 (9) (2022) 4387–4396.

[50] P. L. Woodworth, A. Melet, M. Marcos, R. D. Ray, G. Wöppelmann, Y. N. Sasaki, M. Cirano, A. Hibbert, J. M. Huthnance, S. Monserrat, et al., Forcing factors affecting sea level changes at the coast, Surveys in Geophysics 40 (6) (2019) 1351–1397.

[51] I. D. Haigh, M. Marcos, S. A. Talke, P. L. Woodworth, J. R. Hunter, B. S. Hague, A. Arns, E. Bradshaw, P. Thompson, Gesla version 3: A major update to the global higher-frequency sea-level dataset, Geoscience Data Journal 10 (3) (2023) 293–314.

[52] M. Hart-Davis, D. Dettmering, F. Seitz, Ticon-3: Tidal constants based on gesla-3 sea-level records from globally distributed tide gauges including gauge type information (data), Deutsches Geodätisches Forschungsinstitut, München.

[53] S. P. Neeck, E. J. Lindstrom, P. V. Vaze, L.-L. Fu, Surface water and ocean topography (swot) mission, in: Sensors, Systems, and Next-Generation Satellites XVI, Vol. 8533, SPIE, 2012, pp. 111–120.

[54] CNES, Fes2022 (Finite Element Solution) Ocean Tide (Version 2022), the FES2022 Tide product was funded by CNES, produced by LEGOS, NOVELTIS, and CLS, and made freely available by AVISO. (2024).

[55] R. Ray, G. Egbert, S. Kelly, Barotropic and baroclinic tide models for and from swot.

[56] B. K. Arbic, F. Lyard, A. Ponte, R. D. Ray, J. G. Richman, J. F. Shriver, E. Zaron, Z. Zhao, Tides and the SWOT mission: Transition from science definition team to science team.

[57] M. Hart-Davis, O. Andersen, R. Ray, E. Zaron, C. Schwatke, R. Arildsen, D. Dettmering, K. Nielsen, Tides in complex coastal regions: Early case studies from wide-swath swot measurements, Geophysical Research Letters 51 (20) (2024) e2024GL109983.

[58] D. Lichtman, C. Banks, F. Calafat, C. Gommenginger, P. Bell, Towards validation of SWOT in the coastal zone: a radar altimetry and water level gauge case study in the Bristol Channel and Severn River-Estuary system, in: 2023 Ocean Surface Topography Science Team Meeting, 2023, p. 123.

[59] P. Schaeffer, M.-I. Pujol, P. Veillard, Y. Faugere, Q. Dagneaux, G. Dibarboure, N. Picot, The cnes cls 2022 mean sea surface: Short wavelength improvements from cryosat-2 and saral/altika high-sampled altimeter data, Remote Sensing 15 (11) (2023) 2910.

[60] R. E. Thomson, W. J. Emery, Data analysis methods in physical oceanography, Newnes, 2014.

[61] C. Le Provost, Ocean tides, in: International geophysics, Vol. 69, Elsevier, 2001, pp. 267–303.

[62] I. J. Myung, Tutorial on maximum likelihood estimation, Journal of mathematical Psychology 47 (1) (2003) 90–100.

[63] S. Innocenti, P. Matte, V. Fortin, N. Bernier, Analytical and residual bootstrap methods for parameter uncertainty assessment in tidal analysis with temporally correlated noise, Journal of Atmospheric and Oceanic Technology 39 (10) (2022) 1457–1481.

[64] P. W. Holland, R. E. Welsch, Robust regression using iteratively reweighted least-squares, Communications in Statistics-theory and Methods 6 (9) (1977) 813–827.

[65] M. Mueller, J. Y. Cherniawsky, M. G. Foreman, J.-S. von Storch, Seasonal variation of the m 2 tide, Ocean Dynamics 64 (2014) 159–177.

[66] C. Wunsch, Discrete inverse and state estimation problems: with geophysical fluid applications, Cambridge University Press, 2006.

[67] L. Kachelein, Bayesian Harmonic Analysis of Tidal and Wind-Driven Currents in the California Current System, University of California, San Diego, 2023.

[68] S. J. Roberts, W. D. Penny, Variational bayes for generalized autoregressive models, IEEE Transactions on Signal Processing 50 (9) (2002) 2245–2257.

[69] C. M. Bishop, N. M. Nasrabadi, Pattern recognition and machine learning, Vol. 4, Springer, 2006.

[70] J. J. O. Ruanaidh, W. J. Fitzgerald, Numerical Bayesian methods applied to signal processing, Springer Science & Business Media, 2012.

[71] S. Roberts, A. McQuillan, S. Reece, S. Aigrain, Astrophysically robust systematics removal using variational inference: application to the first month of kepler data, Monthly Notices of the Royal Astronomical Society 435 (4) (2013) 3639–

3653.

[72] C. W. Fox, S. J. Roberts, A tutorial on variational bayesian inference, Artificial intelligence review 38 (2012) 85–95.

[73] W. Penny, S. Roberts, Bayesian multivariate autoregressive models with structured priors, IEE Proceedings-Vision, Image and Signal Processing 149 (1) (2002) 33–41.

[74] A. Dempster, N. Laird, D. Rubin, Maximum Likelihood from Incomplete Data via the EM Algorithm, J. Roy. Stat. Soc. 39 (1) (1977) 1–38.

[75] J. Salvatier, T. V. Wiecki, C. Fonnesbeck, Probabilistic programming in python using pymc3, PeerJ Computer Science 2 (2016) e55.

[76] A. Caticha, R. Preuss, Maximum entropy and bayesian data analysis: Entropic prior distributions, Physical Review E—Statistical, Nonlinear, and Soft Matter Physics 70 (4) (2004) 046127.

[77] H. Attias, A variational baysian framework for graphical models, Advances in neural information processing systems 12.

[78] Y. Wang, A. Kucukelbir, D. M. Blei, Robust probabilistic modeling with bayesian data reweighting, in: International Conference on Machine Learning, PMLR, 2017, pp. 3646–3655.

[79] E. Schrama, R. Ray, A preliminary tidal analysis of topex/poseidon altimetry, Journal of Geophysical Research: Oceans 99 (C12) (1994) 24799–24808.

[80] P. Matte, N. Bernier, J.-M. Fiset, V. Fortin, Y. Secretan, Constrained harmonic analysis: A new method for mapping estuarine tides based on in-situ and remotely-sensed data, in: Poster presented at the 2018 Ocean Sciences Meeting, 2018, pp. 11–16.

[81] L. Anselin, Spatial econometrics: methods and models, Vol. 4, Springer Science & Business Media, 2013.

[82] D. L. Donoho, P. J. Huber, The notion of breakdown point, A festschrift for Erich L. Lehmann 157184.

[83] P. Rousseeuw, A. Leroy, Robust regression and outlier detection (1987).

[84] J. Williams, K. Horsburgh, Evaluation and comparison of the operational Bristol Channel Model storm surge suite, Research and Consultancy Report 38, National Oceanography Centre, Southampton (2013).

[85] T. A. A. Adcock, S. Draper, G. T. Houlsby, A. G. Borthwick, S. Serhadl Serhadlıoğlu, The available power from tidal stream turbines in the Pentland Firth, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 469 (2157) (2013) 20130072.

[86] M. G. Hart-Davis, D. Dettmering, R. Sulzbach, M. Thomas, C. Schwatke, F. Seitz, Regional evaluation of minor tidal constituents for improved estimation of ocean tides, Remote Sensing 13 (16) (2021) 3310.

[87] G. Piccioni, D. Dettmering, C. Schwatke, M. Passaro, F. Seitz, Design and regional assessment of an empirical tidal model based on fes2014 and coastal altimetry, Advances in Space Research 68 (2) (2021) 1013–1022.

[88] Y. Li, G. Peng, P. Chen, K. Chen, R. Li, Z. Song, Harmonic analysis of short-term tidal level prediction model for tidal reaches, Arabian Journal of Geosciences 15 (6) (2022) 473.

[89] H. Pan, T. Xu, Z. Wei, Improved tidal estimates from short water level records via the modified harmonic analysis model, Ocean Modelling 189 (2024) 102372.

[90] N. S. Diffenbaugh, D. Singh, J. S. Mankin, D. E. Horton, D. L. Swain, D. Touma, A. Charland, Y. Liu, M. Haugen, M. Tsiang, et al., Quantifying the influence of global warming on unprecedented extreme climate events, Proceedings of the National Academy of Sciences 114 (19) (2017) 4881–4886.

[91] J. E. Overland, Arctic climate extremes, Atmosphere 13 (10) (2022) 1670.

[92] L. Landrum, M. M. Holland, Extremes become routine in an emerging new arctic, Nature Climate Change 10 (12) (2020) 1108–1115.

[93] A. Carrillo-Ponce, S. Heimann, G. M. Petersen, T. R. Walter, S. Cesca, T. Dahm, The 16 september 2023 greenland megatsunami: Analysis and modeling of the source and a weeklong, monochromatic seismic signal, The Seismic Record 4 (3) (2024) 172–183.

[94] K. Svennevig, S. P. Hicks, T. Forbriger, T. Lecocq, R. Widmer-Schnidrig, A. Mangeney, C. Hibert, N. J. Korsgaard, A. Lucas, C. Satriano, et al., A rockslide-generated tsunami in a greenland fjord rang earth for 9 days, Science 385 (6714) (2024) 1196–1205.

[95] A. B. Rabinovich, Seiches and harbor oscillations, in: Handbook of coastal and ocean engineering, World Scientific, 2010, pp. 193–236.

[96] J. M. Amundson, J. F. Clinton, M. Fahnestock, M. Truffer, M. P. Lüthi, R. J. Motyka, Observing calving-generated ocean waves with coastal broadband seismometers, jakobshavn isbræ, greenland, Annals of Glaciology 53 (60) (2012) 79–84.

[97] J. Macías, C. Escalante, M. J. Castro, Multilayer-hysea model validation for landslide-generated tsunamis–part 1: Rigid slides, Natural Hazards and Earth System Sciences 21 (2) (2021) 775–789.

[98] D. B. Chelton, J. C. Ries, B. J. Haines, L.-L. Fu, P. S. Callahan, Satellite altimetry, in: International geophysics, Vol. 69, Elsevier, 2001, pp. 1–ii.

[99] R. Fjørtoft, J.-M. Gaudin, N. Pourthié, J.-C. Lalaurie, A. Mallet, J.-F. Nouvel, J. Martinot-Lagarde, H. Oriot, P. Borderies, C. Ruiz, et al., Karin on swot: Characteristics of near-nadir ka-band interferometric sar imagery, IEEE Transactions on Geoscience and Remote Sensing 52 (4) (2013) 2172–2185.

[100] L.-L. Fu, T. Pavelsky, J.-F. Cretaux, R. Morrow, J. T. Farrar, P. Vaze, P. Sengenes, N. Vinogradova-Shiffer, A. Sylvestre-Baron, N. Picot, et al., The surface water and ocean topography mission: A breakthrough in radar remote sensing of the ocean and land surface water, Geophysical Research Letters 51 (4) (2024) e2023GL107652.

[101] M. E. Pasyanos, T. G. Masters, G. Laske, Z. Ma, Litho1. 0: An updated crust and lithospheric model of the earth, Journal of Geophysical Research: Solid Earth 119 (3) (2014) 2153–2173.

[102] A. Gelman, B. Goodrich, J. Gabry, A. Vehtari, R-squared for bayesian regression models, The American Statistician.

[103] M. Caceres, A. Valle-Levinson, L. Atkinson, Observations of cross-channel structure of flow in an energetic tidal channel, Journal of Geophysical Research: Oceans 108 (C4).

[104] F. Cottier, F. Nilsen, R. Skogseth, V. Tverberg, J. Skarðhamar, H. Svendsen, Arctic fjords: a re-

view of the oceanographic environment and dominant physical processes, Geological Society, London, Special Publications 344 (1) (2010) 35–50.

[105] T. Monahan, T. Tang, S. Roberts, T. A. A. Adcock, Tidal corrections from and for SWOT using a spatially coherent variational Bayesian harmonic analysis, Journal of Geophysical Research: Oceans 130 (3) (2025) e2024JC021533.

[106] J. F. Price, R. A. Weller, R. R. Schudlich, Wind-driven ocean currents and ekman transport, Science 238 (4833) (1987) 1534–1538.

[107] D. B. Chelton, R. A. DeSzoeke, M. G. Schlax, K. El Naggar, N. Siwertz, Geographical variability of the first baroclinic rossby radius of deformation, Journal of Physical Oceanography 28 (3) (1998) 433–460.

[108] A. Nurser, S. Bacon, The rossby radius in the arctic ocean, Ocean Science 10 (6) (2014) 967–975.

[109] M. Rantanen, A. Y. Karpechko, A. Lipponen, K. Nordling, O. Hyvärinen, K. Ruosteenoja, T. Vihma, A. Laaksonen, The arctic has warmed nearly four times faster than the globe since 1979, Communications earth & environment 3 (1) (2022) 168.

[110] L. Krischer, T. Megies, R. Barsch, M. Beyreuther, T. Lecocq, C. Caudron, J. Wassermann, Obspy: A bridge for seismology into the scientific python ecosystem, Computational Science & Discovery 8 (1) (2015) 014003.

[111] T. Monahan, Observations of the seiche that shook the world (replication materials) (Apr. 2025). doi:10.5281/zenodo.15166491.
URL `https://doi.org/10.5281/zenodo.15166491`

[112] Z. Xu, T. D. Mikesell, Estimation of resolution and covariance of ambient seismic source distributions: Full waveform inversion and matched field processing, Journal of Geophysical Research: Solid Earth 127 (6) (2022) e2022JB024374.

[113] R. SATO, Seismic waves in the near field, Journal of Physics of the Earth 20 (4) (1972) 357–375.

[114] W. Boone, S. Rysgaard, E. Frandsen, R. Develter, F. M. Institute, A. University, Greenland integrated observatory - ctd & atmospheric station dickson fjord - 2023, `https://doi.org/10.14284/637` (2023).

[115] I. Bij de Vaate, A. Vasulkar, D. Slobbe, M. Verlaan, The influence of Arctic landfast ice on seasonal modulation of the M2 tide, Journal of Geophysical Research: Oceans 126 (5) (2021) e2020JC016630.

[116] B. Cushman-Roisin, A. J. Willmott, N. R. Biggs, Influence of stratification on decaying surface seiche modes, Continental Shelf Research 25 (2) (2005) 227–242.

[117] Surface Water Ocean Topography (SWOT), SWOT Level 2 Water Mask Pixel Cloud Auxiliary Data Product, Version C, ver. C. PO.DAAC, CA, USA. Dataset accessed 2025-04-07 (2024). doi:10.5067/SWOT-PIXCVEC-2.0.
URL `https://doi.org/10.5067/SWOT-PIXCVEC-2.0`

[118] P. Matte, Y. Secretan, J. Morin, Temporal and spatial variability of tidal-fluvial dynamics in the st. lawrence fluvial estuary: An application of nonstationary tidal harmonic analysis, Journal of Geophysical Research: Oceans 119 (9) (2014) 5724–5744.

[119] J. Williams, M. Irazoqui Apecechea, A. Saulter, K. J. Horsburgh, Storm surge forecasting: quantifying errors arising from the double-counting of radiational tides, Ocean Sci Discuss 2018 (2018) 1–21.

[120] A. T. Doodson, Oceanic tides, in: Advances in geophysics, Vol. 5, Elsevier, 1958, pp. 117–152.

[121] S. L. Brunton, B. R. Noack, P. Koumoutsakos, Machine learning for fluid mechanics, Annual review of fluid mechanics 52 (1) (2020) 477–508.

[122] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, et al., Learning skillful medium-range global weather forecasting, Science 382 (6677) (2023) 1416–1421.

[123] F. Kratzert, D. Klotz, G. Shalev, G. Klambauer, S. Hochreiter, G. Nearing, Towards learning universal, regional, and local hydrological behaviors via machine learning applied to large-sample datasets, Hydrology and Earth System Sciences 23 (12) (2019) 5089–5110.

[124] D. Sarkar, M. Osborne, T. Adcock, A machine learning approach to the prediction of tidal currents., in: ISOPE International Ocean and Polar Engineering Conference, ISOPE, 2016, pp. ISOPE–I.

[125] D. Sarkar, M. A. Osborne, T. A. A. Adcock, Prediction of tidal currents using Bayesian machine learning, Ocean Engineering 158 (2018) 221–231.

[126] M. M. De Oliveira, N. F. F. Ebecken, J. L. F. De Oliveira, I. de Azevedo Santos, Neural network model to predict a storm surge, Journal of applied Meteorology and Climatology 48 (1) (2009) 143–155.

[127] K. Chen, C. Kuang, L. Wang, K. Chen, X. Han, J. Fan, Storm surge prediction based on long short-term memory neural network in the east china sea, Applied Sciences 12 (1) (2021) 181.

[128] Y. Qin, C. Su, D. Chu, J. Zhang, J. Song, A review of application of machine learning in storm surge problems, Journal of Marine Science and Engineering 11 (9) (2023) 1729.

[129] D. A. Jay, T. Kukulka, Revising the paradigm of tidal analysis–the uses of non-stationary data, Ocean Dynamics 53 (2003) 110–125.

[130] H. Pan, X. Lv, Y. Wang, P. Matte, H. Chen, G. Jin, Exploration of tidal-fluvial interaction in the columbia river estuary using s_tide, Journal of Geophysical Research: Oceans 123 (9) (2018) 6598–6619.

[131] G. W. Groves, R. W. Reynolds, An orthogonalized convolution method of tide prediction, Journal of Geophysical Research 80 (30) (1975) 4131–4138.

[132] J. Donini, B. Hollebone, A. Lever, The derivation and application of normalized spherical harmonic hamiltonians, Progress in Inorganic Chemistry 22 (1977) 225–261.

[133] J. Wray, G. G. Green, Calculation of the Volterra kernels of non-linear dynamic systems using an artificial neural network, Biological cybernetics 71 (3) (1994) 187–195.

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

[135] V. Z. Marmarelis, X. Zhao, Volterra models and three-layer perceptrons, IEEE Transactions on neural networks 8 (6) (1997) 1421–1433.

[136] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, et al., Wide & deep learning for recommender systems, in: Proceedings of the 1st workshop on deep learning for recommender systems, 2016, pp. 7–10.

[137] K. Judd, A. Mees, Embedding as a modeling problem, Physica D: Nonlinear Phenomena 120 (3-4) (1998) 273–286.

[138] E. Tan, S. Algar, D. Corrêa, M. Small, T. Stemler, D. Walker, Selecting embedding delays: An overview of embedding techniques and a new method using persistent homology, Chaos: An Interdisciplinary Journal of Nonlinear Science 33 (3).

[139] A. Fowler, G. Kember, Delay recognition in chaotic time series, Physics Letters A 175 (6) (1993) 402–408.

[140] J. Rala Cordeiro, A. Raimundo, O. Postolache, P. Sebastião, Neural architecture search for 1d cnns—different approaches tests and measurements, Sensors 21 (23) (2021) 7990.

[141] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, A. Talwalkar, Hyperband: A novel bandit-based approach to hyperparameter optimization, The journal of machine learning research 18 (1) (2017) 6765–6816.

[142] M. Wedler, M. Stender, M. Klein, S. Ehlers, N. Hoffmann, Surface similarity parameter: A new machine learning loss metric for oscillatory spatio-temporal data, Neural Networks 156 (2022) 123–134.

[143] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980.

[144] S. M. Lundberg, B. Nair, M. S. Vavilala, M. Horibe, M. J. Eisses, T. Adams, D. E. Liston, D. K.-W. Low, S.-F. Newman, J. Kim, et al., Explainable machine-learning predictions for the prevention of hypoxaemia during surgery, Nature biomedical engineering 2 (10) (2018) 749–760.

[145] A. T. Doodson, Report on Thames floods, no. 7, HM Stationery Office, 1929.

[146] J. Keers, An empirical investigation of interaction between storm surge and astronomical tide on the east coast of great britain, Deutsche Hydrographische Zeitschrift 21 (3) (1968) 118–125.

[147] D. Prandle, J. Wolf, The interaction of surge and tide in the north sea and river thames, Geophysical Journal International 55 (1) (1978) 203–216.

[148] D. Idier, F. Dumas, H. Muller, Tide-surge interaction in the english channel, Natural Hazards and Earth System Sciences 12 (12) (2012) 3709–3718.

[149] L.-L. Fu, G. Pihos, Determining the response of sea level to atmospheric pressure forcing using topex/poseidon data, Journal of Geophysical Research: Oceans 99 (C12) (1994) 24633–24642.

[150] A. F. Hoitink, D. A. Jay, Tidal river dynamics: Implications for deltas, Reviews of Geophysics 54 (1) (2016) 240–272.

[151] J. T. Wells, Tide-dominated estuaries and tidal rivers, in: Developments in sedimentology, Vol. 53, Elsevier, 1995, pp. 179–205.

[152] M. Gan, Y. Chen, S. Pan, J. Li, Z. Zhou, A modified nonstationary tidal harmonic analysis model for the yangtze estuarine tides, Journal of Atmospheric and Oceanic Technology 36 (4) (2019) 513–525.

[153] W. Xie, G. Xu, H. Zhang, C. Dong, Developing a deep learning-based storm surge forecasting model, Ocean Modelling 182 (2023) 102179.

[154] H. Galavi, M. Mirzaei, L. T. Shul, N. Valizadeh, Klang river–level forecasting using ARIMA and ANFIS models, Journal-American Water Works Association 105 (9) (2013) E496–E506.

[155] L. See, S. Openshaw, A hybrid multi-model approach to river level forecasting, Hydrological Sciences Journal 45 (4) (2000) 523–536.

[156] L. See, S. Openshaw, Applying soft computing approaches to river level forecasting, Hydrological Sciences Journal 44 (5) (1999) 763–778.

[157] C. Tebaldi, B. H. Strauss, C. E. Zervas, Modelling sea level rise impacts on storm surges along us coasts, Environmental Research Letters 7 (1) (2012) 014032.

[158] H. R. Moftakhari, A. AghaKouchak, B. F. Sanders, M. Allaire, R. A. Matthew, What is nuisance flooding? defining and monitoring an emerging challenge, Water Resources Research 54 (7) (2018) 4218–4227.

[159] H. R. Moftakhari, A. AghaKouchak, B. F. Sanders, R. A. Matthew, Cumulative hazard: The case of nuisance flooding, Earth's Future 5 (2) (2017) 214–223.

[160] N. Kohno, S. K. Dube, M. Entel, S. Fakhruddin, D. Greenslade, M.-D. Leroux, J. Rhome, N. B. Thuy, Recent progress in storm surge forecasting, Tropical Cyclone Research and Review 7 (2) (2018) 128–139.

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

[162] R. A. Luettich, J. J. Westerink, N. W. Scheffner, et al., Adcirc: an advanced three-dimensional circulation model for shelves, coasts, and estuaries. report 1, theory and methodology of adcirc-2dd1 and adcirc-3dl.

[163] K. Horsburgh, C. Wilson, Tide-surge interaction and its role in the distribution of surge residuals in the north sea, Journal of Geophysical Research: Oceans 112 (C8).

[164] J. L. Rego, C. Li, Nonlinear terms in storm surge predictions: Effect of tide and shelf geometry with case study from hurricane rita, Journal of Geophysical Research: Oceans 115 (C6).

[165] X. Gou, H. Liang, T. Cai, X. Wang, Y. Chen, X. Xia, The impact of coastline and bathymetry changes on the storm tides in zhejiang coasts, Journal of Marine Science and Engineering 11 (9) (2023) 1832.

[166] J. Williams, A. Saulter, C. O'Neill, J. Brown, K. Horsburgh, A reassessment of the uk operational surge forecasting procedure.

[167] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, A. Ramadhan, A. Edelman, Universal differential equations for scientific machine learning, arXiv preprint arXiv:2001.04385.

[168] D. Idier, X. Bertin, P. Thompson, M. D. Pickering, Interactions between mean sea level, tide, surge, waves and flooding: mechanisms and contributions to sea level variations at the coast, Surveys in Geophysics 40 (6) (2019) 1603–1630.

[169] T. Monahan, RTide, `https://github.com/thomasmonahan/RTide` (2024).

[170] L. Blunden, A. Bahaj, Tidal energy resource assessment for tidal stream generators, Proceed-

ings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy 221 (2) (2007) 137–146.

[171] A. Schendel, A. Hildebrandt, N. Goseberg, T. Schlurmann, Processes and evolution of scour around a monopile induced by tidal currents, Coastal Engineering 139 (2018) 65–84.

[172] D. Prandle, The vertical structure of tidal currents, Geophysical & astrophysical fluid dynamics 22 (1-2) (1982) 29–49.

[173] D. Prandle, Tidal currents in shelf seas—their nature and impacts, Progress in Oceanography 40 (1-4) (1997) 245–261.

[174] T. Stiven, S. J. Couch, A. S. Iyer, Assessing the impact of ADCP resolution and sampling rate on tidal current energy project economics, in: OCEANS 2011 IEEE-Spain, IEEE, 2011, pp. 1–10.

[175] D. A. Jay, J. D. Musiak, Internal tidal asymmetry in channel flows: Origins and consequences, Mixing in estuaries and coastal seas 50 (1996) 211–249.

[176] D. A. Jay, E. P. Flinchem, Interaction of fluctuating river flow with a barotropic tide: A demonstration of wavelet tidal analysis methods, Journal of Geophysical Research: Oceans 102 (C3) (1997) 5705–5720.

[177] D. Cartwright, Tidal analysis-A retrospect, in: Developments in Water Science, Vol. 17, Elsevier, 1982, pp. 170–188.

[178] P. L. Woodworth, A note on the nodal tide in sea level records, Journal of Coastal Research 28 (2) (2012) 316–323.

[179] B. D. Zetler, W. H. Munk, The optimum wiggliness of tidal admittances.

[180] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, G. Wetzstein, Implicit neural representations with periodic activation functions, Advances in neural information processing systems 33 (2020) 7462–7473.

[181] R. E. Kass, A. E. Raftery, Bayes factors, Journal of the american statistical association 90 (430) (1995) 773–795.

[182] T. A. A. Adcock, S. Draper, G. T. Houlsby, A. G. Borthwick, S. Serhadlıoğlu, 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) (2014) 854–861.

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

[184] D. Prandle, J. Matthews, The dynamics of nearshore surface currents generated by tides, wind and horizontal density gradients, Continental Shelf Research 10 (7) (1990) 665–681.

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

[186] T. Monahan, T. Tang, S. Roberts, T. A. A. Adcock, 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, no. 838, 2025.

[187] D. Barber, C. M. Bishop, Ensemble learning in bayesian neural networks, Nato ASI Series F Computer and Systems Sciences 168 (1998) 215–238.

[188] C. Frost, Cost reduction pathway of tidal stream energy in the UK and France (2022).

[189] T. A. A. Adcock, S. Draper, R. H. J. Willden, C. R. Vogel, The fluid mechanics of tidal stream energy conversion, Annual Review of Fluid Mechanics 53 (2021) 287–310.

[190] S. P. Neill, M. R. Hashemi, M. J. Lewis, The role of tidal asymmetry in characterizing the tidal energy resource of Orkney, Renewable Energy 68 (2014) 337–350.

[191] M. Beaudin, H. Zareipour, A. Schellenberglabe, W. Rosehart, Energy storage for mitigating the variability of renewable electricity sources: An updated review, Energy for sustainable development 14 (4) (2010) 302–314.

[192] A. Y. Saber, T. Khandelwal, Iot based online load forecasting, in: 2017 Ninth Annual IEEE Green Technologies Conference (GreenTech), IEEE, 2017, pp. 189–194.

[193] A. Kavousi-Fard, W. Su, A combined prognostic model based on machine learning for tidal current prediction, IEEE Transactions on Geoscience and Remote Sensing 55 (6) (2017) 3108–3114.

[194] N. Safari, O. A. Ansari, A. Zare, C. Chung, A novel decomposition-based localized short-term tidal current speed and direction prediction model, in: 2017 IEEE Power & Energy Society General Meeting, IEEE, 2017, pp. 1–5.

[195] D. Sarkar, M. A. Osborne, T. A. A. Adcock, Spatiotemporal prediction of tidal currents using Gaussian processes, Journal of Geophysical Research: Oceans 124 (4) (2019) 2697–2715.

[196] N. Safari, S. M. Mazhari, B. Khorramdel, C. Chung, Tidal current and level uncertainty prediction via adaptive linear programming, IEEE Transactions on Sustainable Energy 10 (2) (2018) 748–758.

[197] A. Zhang, Y. Lin, Y. Sun, H. Yuan, M. Wang, G. Liu, Y. Hu, Tidal current prediction based on fractal theory and improved least squares support vector machine, IET Renewable Power Generation 16 (2) (2022) 389–401.

[198] T. Lee, D. Jeng, Application of artificial neural networks in tide-forecasting, Ocean Engineering 29 (9) (2002) 1003–1022.

[199] B. F. Chen, H. D. Wang, C. C. Chu, Wavelet and artificial neural network analyses of tide forecasting and supplement of tides around Taiwan and South China Sea, Ocean Engineering 34 (16) (2007) 2161–2175.

[200] X. Qiao, F. Guo, R. Zhang, Short-term tidal current prediction based on GA-BP neural network, in: IOP Conference Series: Earth and Environmental Science, Vol. 513, IOP Publishing, 2020, p. 012061.

[201] H. H. H. Aly, M. E. El Hawary, A proposed ANN and FLSM hybrid model for tidal current magnitude and direction forecasting, IEEE Journal of Oceanic Engineering 39 (1) (2013) 26–31.

[202] H. H. Aly, A novel approach for harmonic tidal currents constitutions forecasting using hybrid intelligent models based on clustering methodologies, Renewable Energy 147 (2020) 1554–1564.

[203] P. Remya, R. Kumar, S. Basu, Forecasting tidal currents from tidal levels using genetic algorithm, Ocean engineering 40 (2012) 62–68.

[204] A. Kavousi-Fard, A hybrid accurate model for tidal current prediction, IEEE Transactions on Geoscience and Remote Sensing 55 (1) (2016) 112–118.

[205] P. Qian, B. Feng, X. Liu, D. Zhang, J. Yang,

Y. Ying, C. Liu, Y. Si, Tidal current prediction based on a hybrid machine learning method, Ocean Engineering 260 (2022) 111985.

[206] J. Yin, N. Wang, An online sequential extreme learning machine for tidal prediction based on improved Gath–Geva fuzzy segmentation, Neurocomputing 174 (2016) 85–98.

[207] W. Y. Chang, A literature review of wind forecasting methods, Journal of Power and Energy Engineering 2 (04) (2014) 161.

[208] Z. Hajirahimi, M. Khashei, Hybrid structures in time series modeling and forecasting: A review, Engineering Applications of Artificial Intelligence 86 (2019) 83–106.

[209] Y. K. Wu, J. S. Hong, A literature review of wind forecasting technology in the world, 2007 IEEE Lausanne Power Tech (2007) 504–509.

[210] S. S. Soman, H. Zareipour, O. Malik, P. Mandal, A review of wind power and wind speed forecasting methods with different time horizons, in: North American power symposium 2010, IEEE, 2010, pp. 1–8.

[211] H. Yang, K. Huang, I. King, M. R. Lyu, Localized support vector regression for time series prediction, Neurocomputing 72 (10-12) (2009) 2659–2669.

[212] A. Dolara, F. Grimaccia, S. Leva, M. Mussetta, E. Ogliari, A physical hybrid artificial neural network for short term forecasting of pv plant power output, Energies 8 (2) (2015) 1138–1153.

[213] L. Shen, Z. Shao, Y. Yu, X. Chen, Hybrid approach combining modified gravity model and deep learning for short-term forecasting of metro transit passenger flows, Transportation Research Record 2675 (1) (2021) 25–38.

[214] J. Dong, W. Zeng, L. Wu, J. Huang, T. Gaiser, A. K. Srivastava, Enhancing short-term forecasting of daily precipitation using numerical weather prediction bias correcting with xgboost in different regions of china, Engineering Applications of Artificial Intelligence 117 (2023) 105579.

[215] Y. Cui, L. Jia, W. Fan, Estimation of actual evapotranspiration and its components in an irrigated area by integrating the shuttleworth-wallace and surface temperature-vegetation index schemes using the particle swarm optimization algorithm, Agricultural and Forest Meteorology 307 (2021) 108488.

[216] A. Franco, Tidal Mathematics: The Analysis of Tides., Science 180 (4087) (1973) 729–730.

[217] J. B. Elsner, A. A. Tsonis, Singular spectrum analysis: a new tool in time series analysis, Springer Science & Business Media, 1996.

[218] H. Hassani, D. Thomakos, A review on singular spectrum analysis for economic and financial time series, Statistics and its Interface 3 (3) (2010) 377–397.

[219] K. Afshar, N. Bigdeli, Data analysis and short term load forecasting in Iran electricity market using singular spectral analysis (SSA), Energy 36 (5) (2011) 2620–2627.

[220] C. Bandt, B. Pompe, Permutation entropy: a natural complexity measure for time series, Physical review letters 88 (17) (2002) 174102.

[221] E. Olofsen, J. Sleigh, A. Dahan, Permutation entropy of the electroencephalogram: a measure of anaesthetic drug effect, British journal of anaesthesia 101 (6) (2008) 810–821.

[222] A. A. Pessa, H. V. Ribeiro, ordpy: A Python package for data analysis with permutation entropy and ordinal network methods, Chaos: An Interdisciplinary Journal of Nonlinear Science 31 (6) (2021) 063110.

[223] J. G. Colonna, E. F. Nakamura, Unsupervised selection of the singular spectrum components based on information theory for bioacoustic signal filtering, Digital Signal Processing 82 (2018) 64–79.

[224] H. E. Hurst, Long-term storage capacity of reservoirs, Transactions of the American society of civil engineers 116 (1) (1951) 770–799.

[225] S. Singh, A simple method of forecasting based on fuzzy time series, Applied mathematics and computation 186 (1) (2007) 330–339.

[226] S. M. Chen, Forecasting enrollments based on fuzzy time series, Fuzzy sets and systems 81 (3) (1996) 311–319.

[227] P. C. Silva, P. Lucas, H. Sadaei, F. Guimarães, pyfts: Fuzzy time series for python, Belo Horizonte.

[228] W. Sulandari, S. Subanar, M. H. Lee, P. C. Rodrigues, Time series forecasting using singular spectrum analysis, fuzzy systems and neural networks, MethodsX 7 (2020) 101015.

[229] D. Niyogi, A novel method combines moving fronts, data decomposition and deep learning to forecast intricate time series, arXiv preprint arXiv:2303.06394.

[230] E. J. Kubatko, J. J. Westerink, C. Dawson, hp discontinuous Galerkin methods for advection dominated problems in shallow water flow, Computer Methods in Applied Mechanics and Engineering 196 (1-3) (2006) 437–451.

[231] T. Wang, Z. Yang, A tidal hydrodynamic model for Cook Inlet, Alaska, to support tidal energy resource characterization, Journal of Marine Science and Engineering 8 (4) (2020) 254.

[232] D. M. Kline, Methods for multi-step time series forecasting neural networks, in: Neural networks in business forecasting, IGI Global, 2004, pp. 226–250.

[233] S. B. Taieb, G. Bontempi, A. F. Atiya, A. Sorjamaa, A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition, Expert systems with applications 39 (8) (2012) 7067–7083.

[234] N. Golyandina, On the choice of parameters in singular spectrum analysis and related subspace-based methods, arXiv preprint arXiv:1005.4374.

[235] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825–2830.

[236] F. Chollet, et al., Keras, `https://keras.io` (2015).

[237] R. A. Choudrey, S. J. Roberts, Variational mixture of bayesian independent component analyzers, Neural computation 15 (1) (2003) 213–252.

[238] I. D. Haigh, M. D. Pickering, J. M. Green, B. K. Arbic, A. Arns, S. Dangendorf, D. F. Hill, K. Horsburgh, T. Howard, D. Idier, et al., The tides they are a-changin': A comprehensive review of past and future nonastronomical changes in tides, their driving mechanisms, and future implications, Reviews of Geophysics 58 (1) (2020) e2018RG000636.

[239] D. F. Hill, Spatial and temporal variability in tidal range: evidence, causes, and effects, Current Climate Change Reports 2 (2016) 232–241.

[240] L. Opel, M. Schindelegger, R. D. Ray, A likely role for stratification in long-term changes of the global ocean tides, Communications Earth & Environment 5 (1) (2024) 261.

[241] M. G. Hart-Davis, A. Moiseev, A. Bonaduce, B. C. Backeberg, J. A. Johannessen, Direct ocean tidal current measurements from space: Enhanced interpretation of sentinel-1 doppler shift signals.

[242] M. Garnelo, J. Schwarz, D. Rosenbaum, F. Viola, D. J. Rezende, S. Eslami, Y. W. Teh, Neural processes, arXiv preprint arXiv:1807.01622.