Learning unresolved coastal dynamics in hydrodynamic models

Thomas Monahan[1*], Jeffrey Polton[2], Silvia Innocenti[3], Pascal Matte[3], Mahmoud Ayyad[4], Krijn Saman[5], Thomas A. A. Adcock[1]

[1*]Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ , United Kingdom.

[2] National Oceanography Centre, Joseph Proudman Building, 6 Brownlow Street, Liverpool, L3 5DA, United Kingdom.

[3]Numerical Modelling and Prediction Research Division, Environment and Climate Change Canada, 801-1550 d'Estimauville ave., Quebec, G1J 0C3, Canada.

[4]Department of Civil, Environmental, and Ocean Engineering, Stevens Institute of Technology, 1 Castle Point Terrace, Hoboken, NJ 07030, USA.

[5]Rijkswaterstaat, Ketensedijk 4, LA Capelle aan den IJssel, 2902, Netherlands.

*Corresponding author(s). E-mail(s): thomas.monahan@eng.ox.ac.uk;
Contributing authors: jelt@noc.ac.uk; Silvia.Innocenti@ec.gc.ca; Pascal.Matte@ec.gc.ca; mayyad@stevens.edu; krijn.saman@rws.nl; thomas.adcock@eng.ox.ac.uk;

Abstract

Coastal hydrodynamic models play a vital role in understanding and predicting flooding, but practical computational constraints and uncertainties in boundary conditions and bathymetry lead to systematic errors in local sea level. We show that much of this error is not random but reflects a stable, site-specific response that can be learned from model output and observations. We develop a time-invariant response operator combining linear memory and bilinear interactions to approximate unresolved shallow-water dynamics. In idealized experiments, the operator recovers overtide generation, tide–surge interaction, and fluvial coupling. The approach yields substantial improvements across regional and

operational models of water level and currents, and the correction substantially reduces errors in coarse-resolution simulations, yielding skill comparable to higher-resolution models. Applied globally to a GTSM reanalysis, it reduces a 0.14 m negative bias in 100-year return levels to 0.02 m across 199 GESLA4 gauges, with mean absolute error reductions of 48%, 26%, and 20% for 10-, 50-, and 100-year return periods. These results show that a significant component of coastal model error, including extremes, stems from tidal processes and is locally learnable, offering a practical way to improve skill without modifying underlying models or increasing computational cost.

Keywords: coastal extremes, storm surge, bias correction

1 Introduction

Accurate characterisation of coastal processes, particularly their extremes, is critical for coastal risk assessment worldwide [1, 2]. Physics-based hydrodynamic models, ranging from global tide–surge reanalyses to high-resolution regional and estuarine models, are the principal tool for this purpose [3]. However, these models depend on the accuracy of external forcing, bathymetry, grid resolution, and parameterisations, and must balance fidelity against computational cost. As a result, they routinely exhibit errors that compromise practical applications. For instance, we show in this paper that global coastal sea-level models carry a systematic negative bias of around 0.14 m in 100-year return level estimates even after accounting for mean offsets.

The sources of model error in coastal hydrodynamic models are, in principle, well understood [4]. Dynamical models offer a physically consistent representation of large-scale forcing, but their skill at any given location depends on bathymetric accuracy, grid resolution, and the fidelity with which nonlinear processes are represented [5]. Many of these processes are inherently local, as they depend on channel geometry, depth, and forcing combinations that vary sharply over scales often below those of global or regional model grids [6]. This leads to systematic, site-specific errors that are difficult to resolve through conventional approaches such as increasing model resolution or data assimilation, both of which incur substantial computational cost.

An alternative approach is to learn these structured errors directly from data. Prior work has shown that shallow-water flows can be represented by time-invariant response functions relating forcing to observations [7–9]. This 'Response' approach provides a physics-informed but fully empirical framework that represents local dynamics without explicitly resolving them in the governing equations.

In the present work, we combine the strengths of dynamical models and response-based approaches. Rather than replacing the dynamical model, we learn a time-invariant operator that maps the full modelled signal to local observations, thereby correcting structured model errors while retaining physically consistent large-scale forcing. Building on recent work on linear lag-regressions relating water-level signals from adjacent tide gauges [10], we introduce a bilinear response operator that captures the leading-order nonlinear processes responsible for systematic errors in coastal

hydrodynamic models. The operator is interpretable through its frequency-domain transfer function, and can be learned from only modest observational data. Importantly, the time-invariance of the learned weights allows the learned operator to be applied using model output alone. This makes the operator practically deployable across the global tide-gauge network, and can make use of gauges which are no longer online. This contrasts with recent efforts to apply machine learning to storm surge forecast bias correction [11–13]. While such approaches can deliver strong predictive performance, they are often less interpretable, provide limited insight into the underlying dynamics, and typically rely on concurrent data for autoregressive forecasting, which can limit their applicability in data-sparse settings.

Using the response framework, the contributions of this work are threefold. First, using controlled synthetic experiments, we demonstrate that the proposed bilinear response operator can represent a range of nonlinear shallow-water processes—including overtide generation, tide–surge interaction, and fluvial coupling—and that the associated model errors are largely structured, site-specific, and learnable from short observational records. Second, we demonstrate the utility of these response operators for postprocessing a range of operational models, highlighting the generality of the approach across different model classes, target variables, and modelling deficiencies. Finally, when applying the framework at a global scale, our results show that the method improves both average performance and the representation of extreme water levels, with direct relevance for coastal flood-risk assessment.

2 A time-invariant framework for postprocessing shallow-water models

We treat postprocessing as the problem of learning a stable local response that maps hydrodynamic model output to the observed signal. Let $\hat{\zeta}(t)$ denote the dynamical model prediction and $\zeta(t)$ the corresponding observation, where $\zeta(t)$ may represent either total sea level or currents. We seek a time-invariant operator R such that

$$\zeta(t) \approx R[\hat{\zeta}(t)]. \tag{1}$$

The operator is learned from periods in which model output and observations are both available, but once estimated, it can be applied using model output alone.

The simplest choice for R is a constant bias correction,

$$\zeta(t) \approx b + \hat{\zeta}(t), \tag{2}$$

where b is a learned offset. Such a correction can remove a mean mismatch, but it cannot account for phase errors, delayed responses, or local memory effects. To represent these processes, we instead model R as a finite-memory operator acting on lagged values of the dynamical model output.

The first-order approximation is a linear lag-response model,

$$\zeta(t) \approx b + \sum_{\tau=0}^{L^{(1)}} w_1(\tau)\, \hat{\zeta}(t-\tau), \tag{3}$$

where $w_1(\tau)$ is a learned lag kernel and $L^{(1)}$ is the maximum lag considered for the first-order kernel. The kernel is the full collection of lag-dependent weights $\{w_1(0), w_1(1), \ldots, w_1(L^{(1)})\}$, which defines the memory structure of the correction.

In the time domain, $w_1(\tau)$ quantifies the contribution of model output at lag τ to the corrected estimate, allowing the postprocessed solution to adjust timing, damping, and local persistence. In the frequency domain, the same kernel defines a linear transfer function $H_1(\omega)$ such that

$$\zeta(\omega) \approx H_1(\omega)\hat{\zeta}(\omega), \tag{4}$$

where $\zeta(\omega)$ and $\hat{\zeta}(\omega)$ represent the Fourier transforms of $\zeta(t)$ and $\hat{\zeta}(t)$ respectively. $H_1(\omega)$ is complex-valued and encodes both amplitude and phase corrections. This transfer function has the same interpretation as an admittance in classical tidal response analysis [7, 14].

A linear lag model can only modify the amplitude and phase of existing spectral components in $\hat{\zeta}$. It cannot generate new frequencies. This is a serious limitation for shallow-water applications, where nonlinear processes such as bottom friction, advective acceleration, and time-varying effective depth generate overtides and other interaction products. To capture these effects, we augment the linear response with bilinear interactions between lagged model values. The resulting second-order model is

$$\zeta(t) \approx b + \sum_{\tau=0}^{L^{(1)}} w_1(\tau)\, \hat{\zeta}(t-\tau) + \sum_{\tau_1=0}^{L^{(2)}} \sum_{\tau_2=\tau_1}^{L^{(2)}} w_2(\tau_1,\tau_2), \hat{\zeta}(t-\tau_1)\hat{\zeta}(t-\tau_2), \tag{5}$$

where $w_2(\tau_1, \tau_2)$ is a bilinear kernel and we restrict $\tau_1 \leq \tau_2$ to enforce kernel symmetry. In the time domain, w_2 describes how pairs of lagged model values combine to produce a correction that depends on past state and amplitude. In the frequency domain, this term introduces cross-frequency coupling, allowing energy at one set of frequencies to influence another.

The bilinear term provides the lowest-order mechanism by which the postprocessing operator can represent weakly nonlinear shallow-water dynamics. Products of lagged signals generate sum and difference frequencies, as well as higher harmonics, so the bilinear model can recover leading-order nonlinear behaviour that is absent from a purely linear correction. In this sense, the postprocessing operator is not simply adjusting bulk bias or phase: it is learning a compact, site-specific approximation to unresolved local dynamics.

More generally, the corrected signal may be written as a truncated Volterra series,

$$\zeta(t) \approx b + \sum_{n=1}^{P} \sum_{\tau_1,\ldots,\tau_n=0}^{L^{(n)}} w_n(\tau_1, \ldots, \tau_n) \prod_{j=1}^{n} \hat{\zeta}(t-\tau_j), \tag{6}$$

where P denotes the truncation order of the series, and each kernel w_n is time-invariant and represents nonlinear memory effects of order n. Each kernel w_n acts over its own memory window $L^{(n)}$, with $L^{(1)}$ and $L^{(2)}$ recovering the linear and bilinear forms above. In the present study, we focus on the linear and bilinear terms, which provide the simplest interpretable extension beyond bias correction while remaining tractable to estimate from observational records.

Lag selection

The frequency characteristics of the linear and bilinear operators are determined by the selection of the maximum lags $L^{(1)}$ and $L^{(2)}$ and the lag spacing. In the frequency domain, the linear and bilinear operators can be thought of as filters. The maximum lags dictate the sharpness of the spectral filter and the spacing controls the rate at which the corresponding admittance can vary with frequency. The theoretical properties of these filters and an evaluation of how they behave with different amounts of data are provided in the Supplementary Information in Figures S1-5.

In principle, the approach allows for a distinct memory length $L^{(n)}$ for each kernel, which is both physically motivated and in some locations practically important. Different physical processes act over different timescales: tidal nonlinearities such as overtides and compound tides arise from interactions between constituents whose periods differ by at most a few days, so the bilinear kernel w_2 needs only to extend over a tidal memory window. In contrast, slower forcings, such as seasonal variations in river discharge, stratification, or mean sea level, enter the response additively and are naturally captured by extending the linear memory $L^{(1)}$ without introducing corresponding long-lag interaction terms. Restricting higher-order kernels to physically meaningful and statistically relevant windows also controls parameter count which would otherwise scale as $\binom{L^{(n)}+n}{n}$ (or $(L^{(n)}+1)^n$ without symmetry constraints) and quickly become ill-posed to estimate from finite records. Selection of bespoke values for maximum lags requires knowledge of the processes present at a given location. For simplicity, we here adopt a maximum linear lag of $L^{(1)} = 24$ hours , a maximum bilinear lag of $L^{(2)} = 24$ hours, and a lag spacing of $\tau = 1$ hour. This choice resolves all major tidal constituents (Nyquist 0.5 cph) while limiting parameter growth and reflects the assumption that the model-data mismatch comes from interactions of tidal processes. This selection also explicitly neglects low-frequency interactions, as our primary interest in this work is to improve the representation of high-frequency processes. However, as will be discussed, evaluating longer period processes like fluvial effects warrants future work and can be achieved with different kernel choices. For all the presented applications, we adopt an identical linear lag configuration. The linear and bilinear filters are therefore composed of 25 and 325 terms, respectively.

Weight estimation

Although the bilinear model represents nonlinear interactions in the signal components, it remains linear in its coefficients. This means the unknown weights can be estimated using standard regression tools once the lagged linear and bilinear terms have been assembled into a design matrix. Given a sample of N observations at times $t_1, t_2 \dots t_N$ as $\boldsymbol{\zeta} = [\zeta(t_1), \zeta(t_2), \dots, \zeta(t_N)]^\mathsf{T}$, the postprocessing problem can be

expressed as

$$\boldsymbol{\zeta} = \mathbf{w}^{\mathsf{T}} X + \boldsymbol{\epsilon}, \tag{7}$$

where $\mathbf{w}$ contains the coefficients associated with the linear and bilinear kernels, X is the design matrix whose rows are composed of lagged model values and their bilinear products, and $\boldsymbol{\epsilon}$ is a residual term.

In this study, we estimated $\mathbf{w}$ using a variational Bayesian estimator with Automatic Relevance Determination (ARD) priors [15, 16]. The ARD prior regularizes the solution by shrinking weakly supported coefficients toward zero, which improves robustness to noise and yields a sparser, more interpretable operator. Full details of the estimator are provided in Section 8. Comparisons with other common estimators are provided in the Supplementary Information Section 5.

3 Bilinear response operators recover nonlinear shallow-water interactions

To show that the bilinear response operator can capture unmodeled coastal dynamics, we simulated data with known dynamics to explicitly test the approaches capacity to represent various nonlinear processes. Specifically, we use a 1-D model to simulate tide and surge propagation on a shelf and the additional fluvial and other effects of estuarine environments (See Section 8). Six scenarios are considered, which selectively evaluate different nonlinear mechanisms and features that may be missed in dynamical models and are described in Table 1. To simulate the postprocessing task, we fix a baseline scenario (S0) with all depth-varying nonlinear mechanisms turned off. This is understood to be the numerical model we wish to post-process. Synthetic observational data is taken from the same point in the simulated domain from one of the six different scenarios (S1-S6). Figure 1 shows the result of learning the linear and bilinear postprocessing frameworks to each scenario trained on a 365-day-long hourly time series and tested on a subsequent 60-day time series. Details of the metrics used to evaluate the approach are provided in the Methods.

Tide-surge interaction

We first consider the role of depth-dependent transport in the continuity equation 19. This mechanism generally dominates tide-surge interaction. In scenario S1 (see Table 1, the nonlinear transport term $H \cdot u = (h + \zeta)u$ couples tidal velocity to the surge-elevated water depth, generating energy at sum and difference frequencies of the tide and surge. This is clearly visible in the Weighted Quadratic Transfer Function (WQTF) for S1 (see Figure 1), which shows bright bands concentrated along $f_1 = f_{M_2}$ and $f_2 = f_{M_2}$, indicating that M_2 is the dominant participant in the active nonlinear interactions. The baseline MAE of 1.311 m is reduced to 0.170 m and 0.150 m by the linear and bilinear schemes, respectively, confirming that a significant fraction of the error is attributable to nonlinear frequency coupling that the bilinear kernel can represent but the linear filter cannot. We also evaluate the impact on model performance at extremes in Figure S14 of the Supplementary Information, which shows the 99th-percentile Brier Score, peak water-level height RMSE, and peak timing MAE for all scenarios. For S1, both linear and bilinear approaches improve the peak timing

and height compared to the baseline case; however, the addition of bilinear terms does not impact the representation of extremes.

Case S2 isolates the depth-dependent momentum source term produced by the wind stress term $\tau/\rho H$ and bottom drag C_d/H in Equation 20, through which a surge-elevated water column dilutes the surface momentum input. This is a weaker nonlinear mechanism: the baseline MAE is only 0.251 m, and the WQTF is correspondingly dark, with only faint structure near the M_2–M_2 interaction. The near-zero bilinear residual (MAE 0.001 m) suggests that this nonlinearity is well-approximated by a low-order bilinear correction. This is further confirmed by the results in Figure S14, which shows that the bilinear operator completely corrects timing errors (MAE = 0 minutes). In contrast, the linear and baseline models have errors of 13 and 44 minutes, respectively.

Case S3 combines both mechanisms. The WQTF recovers the M_2-dominated interaction structure seen in S1, but with greater intensity, consistent with the larger baseline error (1.548 m). The linear and bilinear corrections achieve reductions to 0.258 m and 0.230 m, respectively. The fact that the combined error is not simply additive reflects the nonlinear coupling between the two mechanisms, which the bilinear kernel partially captures through cross-frequency interactions.

Unmodeled features

Approximately 25.7% of coastal regions with depths less than 200 m have been directly surveyed [17], meaning that much of the global coastal bathymetry used in hydrodynamic models is derived from sparse observations and interpolation. As a result, significant uncertainty exists in the representation of nearshore geometry depending on the bathymetric product, which can propagate into model errors. We consider two idealized examples of how such representation errors manifest in cases S4 and S5, corresponding to unmodeled bathymetric structure and unmodeled estuarine width convergence, respectively.

The effect of neglecting bathymetric detail (S4) is primarily a frequency-dependent phase and amplitude error in the propagating tidal wave, consistent with the Green's Law framework of [18]: errors in depth h alter the local wave speed $c \propto (gh)^{1/2}$ and the topographic funnelling of tidal amplitude. This manifests as a broadband but spatially diffuse WQTF, reflecting the absence of any dominant interaction frequency pair — the admittance error is distributed across the tidal spectrum rather than concentrated at specific overtide or compound frequencies. Accordingly, the linear scheme captures the bulk of this error (MAE reduced from 0.931 m to 0.054 m, a 95% reduction), and the bilinear scheme offers only modest further improvement (0.026 m, 97% total reduction). Both linear admittances remain close to unity across tidal frequencies, indicating that the postprocessor acts primarily as a broadband amplitude and phase corrector rather than a frequency-transfer operator. The bilinear operator reduces the height and timing errors to 0.01 m and 0 minutes, respectively, indicating that the impact these unmodeled effects on extremes have can almost completely be learned.

Unmodeled width convergence (S5) presents a qualitatively different challenge. In a convergent estuary, topographic funnelling amplifies the incident tidal wave at a rate governed by the convergence and friction [18]. When this convergence is absent from the model geometry, the baseline MAE rises to 3.655 m — the largest of all scenarios.

The linear admittance plot reveals strong frequency-dependent amplification reaching $|H(f)| \approx 100$ at subdiurnal frequencies $f < 0.05$, indicating that the postprocessor must apply large frequency-selective gains to compensate for the missing topographic funneling. The WQTF shows elevated nonlinear energy concentrated at frequencies near the M_2–M_2 interaction, consistent with the dominance of convergence-driven amplitude modulation at the forcing frequency. The bilinear scheme offers substantial relative improvement over the linear correction (MAE reduction of 39%), as well as a reduction in Peak-height RMSE of 53% (Figure S14), suggesting that a large portion of the error comes from a transfer of energy to low and high frequencies. The remaining error arises from the spatially varying, nonlocal nature of convergence-driven amplification that a point-wise bilinear kernel cannot fully represent.

Fluvial dynamics

Neglecting fluvial interactions (S6), even under constant river outflow, leads to significant discrepancies. River flow enters principally through the momentum equation 20, where it augments nonlinear frictional effects and modifies convective accelerations [18]. The quadratic friction term $C_d u|u|$ directly forces overtides at harmonics of M_2, with the third harmonic M_6 receiving particular energy through cubic velocity terms [19]. Critically, these interactions transfer energy to frequencies absent from the original tidal forcing. A linear postprocessor, which can only re-weight existing spectral components, cannot generate these new frequency components — reflected in the degradation of the linear scheme relative to the baseline (MAE 1.358 m vs 0.954 m). The bilinear scheme, by contrast, reduces the error to 0.122 m by explicitly representing cross-frequency interactions. This manifests not just in the average performance, but also in high waters, where the bilinear approach reduced peak water level height errors from 1 m to 0.19 m, respectively (See Figure S14). Unlike the tide-surge cases S1–S3, the WQTF shows a strongly structured pattern of nonlinear interactions extending above 0.2 cph, consistent with river flow modifying tidal dynamics through strong frictional effects [20]. Here, the steady river outflow gives rise to three new sets of harmonics: (i) low frequencies, (ii) frequencies around existing frequencies, and (iii) higher frequencies. Our filter, by design, does not account for low-frequency coupling, so we neglect discussion of these effects (See Section 5 and 7 for a note on these effects). Inspection of the linear admittance and the WQTF shows a spreading of energy around the semi-diurnal band, reflecting the spread of energy to new harmonics such as L_2, μ_2, and ν_2. Finally, the WQTF indicates that the strongest interactions occur between M_6 and other frequency bands. Notably, this is the only case where an interaction is stronger than M_2 and reflects the triple interactions induced by outflow.

4 Learned response offsets deficits in model resolution

Higher spatial resolution is often treated as the primary route to improved hydrodynamic model skill. Higher resolutions capture finer bathymetric structure, sharper coastal geometry, and enable smaller-scale dynamics to be represented explicitly. In practice, however, these gains come at substantial computational cost and are often prohibitive for long reanalyses, large ensembles, global applications, or time-critical

Table 1 Scenario definitions. $H = h + \eta$ is the total (time-varying) depth; h is the still-water depth. The depth variables $\tilde{H}$ and $\hat{H}$ are the depths in the momentum and continuity equations, respectively. Scenarios S0–S3 use a flat shelf with uniform width and a wall at $x = L$. Scenarios S4–S6 introduce estuarine geometry and/or river discharge; the regression Baseline ζ is always S0 on the flat shelf (except for S6, which uses a tapered estuarine geometry with outflow). Full details of the implementation are given in the Methods.

	Scenario	$\tilde{H}$	$\hat{H}$	$h(x)$	$b(x)$	Q_r
Tide–surge interaction (flat shelf)						
S0	Baseline	h	h	flat (tapered for S6)	1	0 ($Q_r > 0$ for S6)
S1	Depth-dep. transport	H	h	flat	1	0
S2	Depth-dep. momentum	h	H	flat	1	0
S3	Combined interaction	H	H	flat	1	0
Geometry and fluvial forcing (estuary)						
S4	Unmodeled bathymetry	h	h	tapered	1	0
S5	Width convergence	h	h	flat	$b(x)$	0
S6	Fluvial–tidal interaction	H	h	tapered	$b(x)$	$Q_r > 0$

operational models. As a result, depending on the application, simulations are often run at resolutions that leave important local processes unresolved. If these missing processes are systematic rather than random, they may be recoverable through learned postprocessing.

To test this idea, we applied the bilinear postprocessing framework to two configurations of the UK NEMO [21] surge model with different horizontal resolution: AMM15 (1.5 km) [22, 23] and AMM75 (7.5 km). These resolutions reflect the grids used to run the models and therefore capture any change in dynamics which come from this choice. Performance was evaluated at 43 Class A tide gauges from the UK national network over 1993–2015 (Figure 2. The respective model output gridpoints are selected based on gauge proximity and will therefore be order 1.5 / 2 km for the AMM15 and 7.5 / 2 km for AMM75 away from the true gauge. This is a potential source of discrepency between models in addition to the unmodeled dynamics.

At each site, postprocessing weights were trained separately for each year and evaluated on all remaining years, yielding a leave-one-year-out assessment of out-of-sample skill. Postprocessing improves both models, but the gains are largest for the coarser configuration and in dynamically complex coastal settings. In the Bristol Channel, for example, the mean absolute error is reduced by 42.3% for AMM15 and 71.1% for AMM75. The clearest illustration occurs at Hinkley Point, near the mouth of the Severn Estuary, where strong spatial gradients in tidal propagation generate pronounced overtides and waveform distortion [24]. Raw MAE is 0.14 m for AMM15 and 0.40 m for AMM75, but after postprocessing, these fall to 0.13 m and 0.11 m, respectively. Thus, although the higher-resolution model resolves more of the local dynamics directly, the residual skill gap can be closed by learning a site-specific bilinear correction.

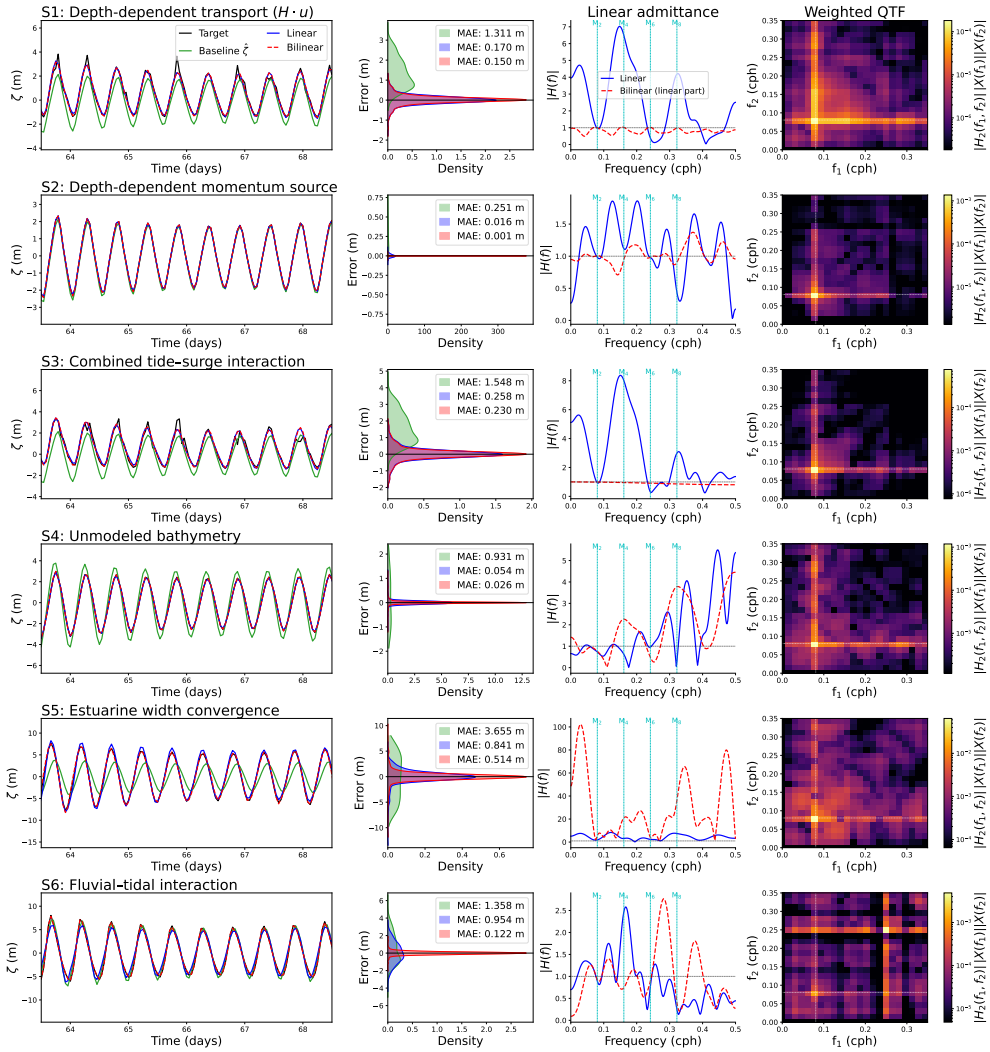

Fig. 1 Evaluation of linear and bilinear postprocessing on synthetic 1-D tide-surge-fluvial model. Each row corresponds to a different experimental configuration with different nonlinear physics enabled. (See Table 1). The baseline model across all cases does not include these physics (described in Table 1). *Columns 1–2*: postprocessed water-level time series and residual error distributions. *Column 3*: linear admittance $|H(f)|$; cyan lines mark M_2, M_4, M_6, M_8. *Column 4*: input-weighted quadratic transfer function $W(f_1, f_2) = |H_2(f_1, f_2)|\,|X(f_1)|\,|X(f_2)|$ (log scale), highlighting energetically active nonlinear frequency interactions. Note that bilinear kernels are symmetric by definition.

This behaviour is consistent across the network. Postprocessing yields a larger relative improvement for the 7.5 km model, yet the postprocessed MAE and 99th-percentile Brier scores of AMM15 and AMM75 are statistically indistinguishable in the

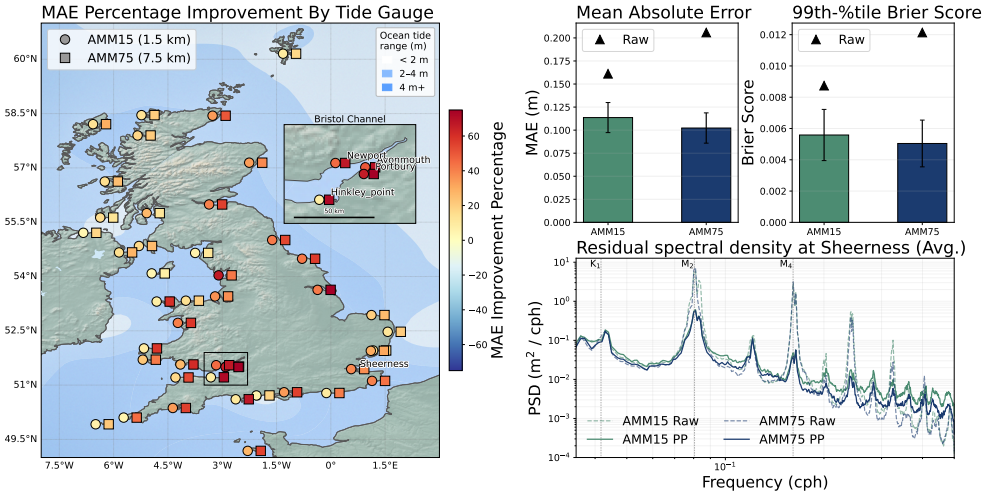

Fig. 2 Comparison of bilinear postprocessing (PP) improvements for two dynamical model configurations with differing grid resolutions (AMM15: 1.5 km; AMM75: 7.5 km). The map shows site-level mean absolute error (MAE) improvement, expressed as a percentage, $100 \times \left(1 - \frac{\mathrm{MAE_{PP}}}{\mathrm{MAE_{raw}}}\right)$, with AMM15 (circles) and AMM75 (squares), Red = postprocessing better. postprocessing weights are computed separately for each individual year and evaluated on all other available years over the period 1993–2020; results are averaged across all such trials. Ocean tidal range is estimated from the FES2022b tidal atlas [25]. Made with Natural Earth. The top-right panels summarise performance across all sites and years in terms of MAE and the 99th percentile Brier score, with raw model values indicated by the black triangles, and standard deviations labelled. The bottom-right panel shows the multi-year mean residual power spectral density before and after postprocessing, with key tidal constituents highlighted. The mean offset over the train period is subtracted from the respective models to allow fair comparison.

multi-site mean (Figure 2). These results suggest that, where in situ observations are available, part of the performance deficit associated with coarse resolution is locally recoverable without refining the underlying grid.

The mechanism of this improvement is apparent in the residual spectra. At Sheerness (Figure 2, bottom right), the raw residual energy of both models is concentrated around M_2 and its higher harmonics, particularly M_4 and M_8, indicating persistent deficiencies in the representation of tidal distortion and nonlinear interactions. Postprocessing strongly suppresses these bands, showing that it is correcting structured tidal error rather than merely reducing broadband variance. It also narrows the spectral cusps around the constituent peaks and lowers the residual spectral floor, consistent with improved representation of both harmonic and non-tidal variability. Unlike a conventional tidal analysis of the residual, which can only remove variability at prescribed tidal frequencies, the bilinear correction acts on the full model time series and can therefore account for interactions between tidal and non-tidal processes.

5 Comparisons with operational models

The synthetic experiments and UK NEMO case study indicate that the learned response operator can recover structured local model error arising from unresolved dynamics and coarse resolution. A key question is whether these gains are specific to a particular model architecture or whether they transfer across models with different governing equations, calibration practices, and target variables. Table 2 addresses this by comparing postprocessing performance across 5 independent simulations from 4 operational forecasting models, a description of which is given in Section 8.

The results show that the framework is not limited to a single coastal sea-level model. Improvements are obtained across multiple systems, including conventional water-level models (DCSM7, Steven's Institute), tidal-current models (MIKE21), and tidal-river applications (Canadian SHOP). This suggests that the method is acting on a broad class of structured hydrodynamic errors rather than exploiting idiosyncrasies of one model configuration. These models are expected to be of the highest quality due to their operational importance. For example, water level forecasts are of direct importance for management and maintenance, as well as for the closure procedures of the storm surge barriers in the Netherlands. The improved MAE performance of DCSM7 of 17.98% for the bilinear operator suggests the approach may benefit such decision making.

Skill gains for the MIKE21 tidal current model stress that the same time-invariant framework extends beyond free-surface elevation to velocity predictions, with even larger gains. The fact that the linear model does marginally better than the bilinear model suggests that the existing model captures the relevant frequency content at the measurement locations; however, this energy is incorrectly distributed. Given there are no additional interactions with other types of forcing under consideration, a linear model is sufficient to represent this variability.

The Steven's Institute Flood Advisory System (SFAS) [26] provides a unique opportunity to evaluate how the proposed approach can be combined with existing postprocessing techniques. The SFAS-TC configuration already includes a dedicated tidal correction (See Section 8), yet even greater improvement is obtained after applying the learned bilinear operator. This indicates that the present approach is not simply reproducing a tidal harmonic adjustment. Rather, it is correcting additional structured errors, including phase-dependent and interaction errors involving non-tidal variability. These are observed by taking the difference between the learned admittances and WQTFS as shown in Supplementary Figure S16. We observe that the TC model shows comparatively less energy at frequency interactions with M2, and instead shows greater energy around interactions between high and low frequency components. Quantitatively, this is further reflected by the fact that the bilinear approach yields a larger improvement in MAE. By injecting additional frequency structure into the model, particularly for compound tides, the tidal correction provides the response operators more structure with which to modulate. Given these results, the method can be viewed as complementary to existing correction strategies rather than as a replacement for them.

The application to operational systems also suggests that the method is useful in environments where nonlinear shallow-water interactions are especially important.

Tidal-river systems are characterised by strong local controls from bathymetry, friction, channel geometry, and flow interaction, all of which contribute to the uncertainty of operational forecasts. The Canadian SHOP model covers a 400 km stretch along the St. Lawrence River and fluvial estuary, encompassing a tide-surge dominated region in the lower estuary and extending to a fluvially dominated region near Montreal. The spatially aggregated percentage improvement shows that the bilinear and linear schemes improve average performance by 14.95% and 12.35% respectively. As expected, the model errors reflect different dynamics throughout the estuary (See Figure S7) with significant differences in postprocessed performance depending on dynamic regime. A map of post-processed performance improvements is also included (Figure S8). A clear shift in dynamics occurs 150 km into the estuary between Portneuf and Batiscan. Across the 5 stations downstream of Batiscan, postprocessing with the bilinear operator yields an average improvement in MAE of 51% (the largest among all operational models). Once the dynamics become fluvially dominated, however, the postprocessing offers no benefit and even degrades model performance (average 6% reduction of MAE). This reflects an important aspect of our model construction; our approach constructs a broadband filter across the tidal spectrum, reflecting an assumption that the relevant dynamics occur in this range. These results suggest the errors in the Canadian SHOP model in these regions sit outside this range and thus warrant modifications to the model construction. Specifically, the residual autocorrelation functions in Figure S7 suggest that in the fluvially dominated region the model is missing dynamics with time-scales greater than 30 days. As such, one would need a maximum lag of at least 30 days and a larger spacing between lags for computational feasibility. Further improvements may be achievable by separating out low and high frequencies through selective filtering – however, such an approach may be prohibitive depending on the data. We thus leave this for future work.

6 Global impacts on extreme sea-level estimation

Global representation of tides and surges presents a challenge for hydrodynamic modeling frameworks due to the computational expense of running models at sufficient resolution and the uncertainties present in coastal geometry products for many regions [17, 27]. Despite this, global models are useful for informing regional flood risk and providing insights into how such risk is evolving. A key question about the time-invariant operator is whether it can improve extreme sea-level estimates. If long term non-stationary processes such as sea-level rise are represented by the hydrodynamic model, this begs the question of whether the time-invariant response weights, which impact high-frequency tidal processes, can be applied decades after they are learned. To evaluate how the approach behaves across a wide range of dynamic regimes, we apply the framework without modification to all GESLA4 gauges [28] with at least 50 years of data and 75% coverage over 1950–2024 and that don't present any detectable (larger than 0.1 m) datum shift over this period. This leaves the set of 199 gauges listed in the Supplementary Information (Table S1) Our analysis is carried out using a leave-one-year-out protocol in which weights trained on a single year are evaluated

Table 2 Model performance comparison for linear and bilinear postprocessing across multiple hydrodynamic models. Raw = raw model output, PP = post-processed. Operators are trained on the 1st year of available data, except for the Stevens SFAS and MIKE21 models, which have only 1 year of data. Reported values are the mean absolute error (MAE) in meters estimated on the remaining, unseen, data before and after postprocessing, along with the mean percentage improvement. DCSM7 = Dutch Operational Surge Forecast System, MIKE21 = MeyGen site operational tidal current forecast, Steven's SFAS = Operational coupled hydrological-coastal forecast system (USA), Canadian SHOP = St. Lawrence Estuary operational forecast system.

Model	Response Operator Type	N_{sites}	MAE (raw)	MAE (PP)	Improvement (%)
DCSM7	Bilinear	3	0.076 m	0.057 m	17.98
DCSM7	Linear	3	0.076 m	0.060 m	14.26
MIKE21 currents	Bilinear	2	0.660 m	0.478 m	26.31
MIKE21 currents	Linear	2	0.660 m	0.459 m	28.73
Stevens SFAS	Bilinear	1	0.061 m	0.059 m	3.13
Stevens SFAS	Linear	1	0.061 m	0.061 m	0.09
Stevens SFAS TC	Bilinear	1	0.057 m	0.053 m	6.73
Stevens SFAS TC	Linear	1	0.057 m	0.055 m	3.21
Canadian SHOP	Bilinear	14	0.197 m	0.187 m	14.95
Canadian SHOP	Linear	14	0.197 m	0.191 m	12.35

on all subsequent years. Average statistics for each site over 10 separate years are reported in Figure 3.

The results show consistent improvement in both mean and high-quantile sea-level skill across the global network compared to GTSM (Figure 3). Given the bilinear operator by construction does not represent mean sea level trends, this only reflects the impact of representing high-frequency tidal processes. Reductions in mean absolute error are widespread, averaging 46% globally, with the largest gains concentrated along dynamically complex coastlines, particularly those with large tidal ranges. The global distribution of Brier score improvements broadly mirrors this pattern, showing a 41% average improvement, confirming that 99th percentile sea-level variability is also better represented after postprocessing. As expected, the benefits of our approach are most significant in areas of large tidal ranges where hydrodynamic models may struggle [29]. postprocessing using default settings, as used here, performs poorly in areas with small astronomical tides such as the Baltic and Mediterranean. A visualization of these statistics as a function of tidal range confirms these observations (Figure S9 in the Supplementary Information).

The most striking result concerns high sea level quantiles. The raw model exhibits a systematic negative bias in 100-year return level estimates of around 0.14 m on average, even after bias removal (see methods). This highlights a global tendency to underestimate extreme events that has direct consequences for flood risk assessment. Postprocessing reduces this mean bias to 0.02 m, nearly an order-of-magnitude correction, while simultaneously improving site-level return level accuracy, reducing the MAE from 0.25m to 0.2m. This suggests that a substantial fraction of the raw model's

extreme-value deficit is not random but reflects structured, learnable errors in the representation of tidal processes. Figure S10 in the Supplementary Information presents an analogous figure using the linear operator. Interestingly, the linear postprocessing scheme increases the model bias across the 199 gauges (mean bias of -0.17 m) whilst yielding a similar reduction in average error (MAE = 0.19 m). This suggests the systematic underestimation of extremes stems from the misrepresentation of nonlinear processes in [30]. Similar results are observed for shorter return periods also – Supplementary Figure S15 shows that the bilinear approach reduces the MAE for 10 and 50 year return periods by 48% and 26% respectively. The increase in improvement percentage for shorter return periods makes sense, as the magnitude of the tidal processes is closer in magnitude to the storm events that produce the extremes. In contrast, for larger return periods, it is expected that the storms contribute relatively more to the total water level and thus errors will be dominated by inaccuracies in the forcing.

The present analysis was restricted to sites with more than 50 years of data in order to perform a rigorous evaluation of the impact of return periods. Our results show that the bilinear response weights learned from one year of data can improve sea-level estimates both on average and in high-quantiles decades into the future. Applying this to the GTSM reanalysis has shown that improving the representation of high-frequency processes can improve return level estimates for flood risk projection; however, this relies on the accurate representation of nonstationary processes like sea-level rise in the hydrodynamic model. If such processes are neglected, the time-invariance of the learned weights no longer holds. These caveats notwithstanding, we identify 2559 sites in the GESLA4 dataset with at least 1-year of data and apply our postprocessing scheme. These weights are made available.

7 Discussion

The results establish that structured, physically interpretable biases in coastal models are widespread, mechanistically identifiable, and correctable from minimal observational data. The synthetic experiments demonstrate that a bilinear response operator can recover the leading-order nonlinear processes responsible for these biases and make them interpretable through the frequency-domain transfer functions. This interpretability is not merely diagnostic: it establishes that the operator is learning physically meaningful structure rather than overfitting noise, which gives confidence in its out-of-sample performance–something that remains a significant concern for purely data-driven approaches. The resolution experiments show that, where local skill is the objective, inaccuracies produced by lower grid resolutions can be learned through postprocessing — two models differing by a factor of five in grid spacing are made statistically indistinguishable after correction. This suggests that coarser configurations may be feasible for long reanalyses and large ensemble applications where pointwise skill at key locations is the evaluation criterion. Validation across seven independent operational models, including both total water level and tidal currents, suggests that these gains are not specific to any particular model architecture or error structure.

At a global scale, these properties translate into a reduction in the systematic underestimate of 100-year return levels from 0.14 m to 0.02 m across the GESLA4

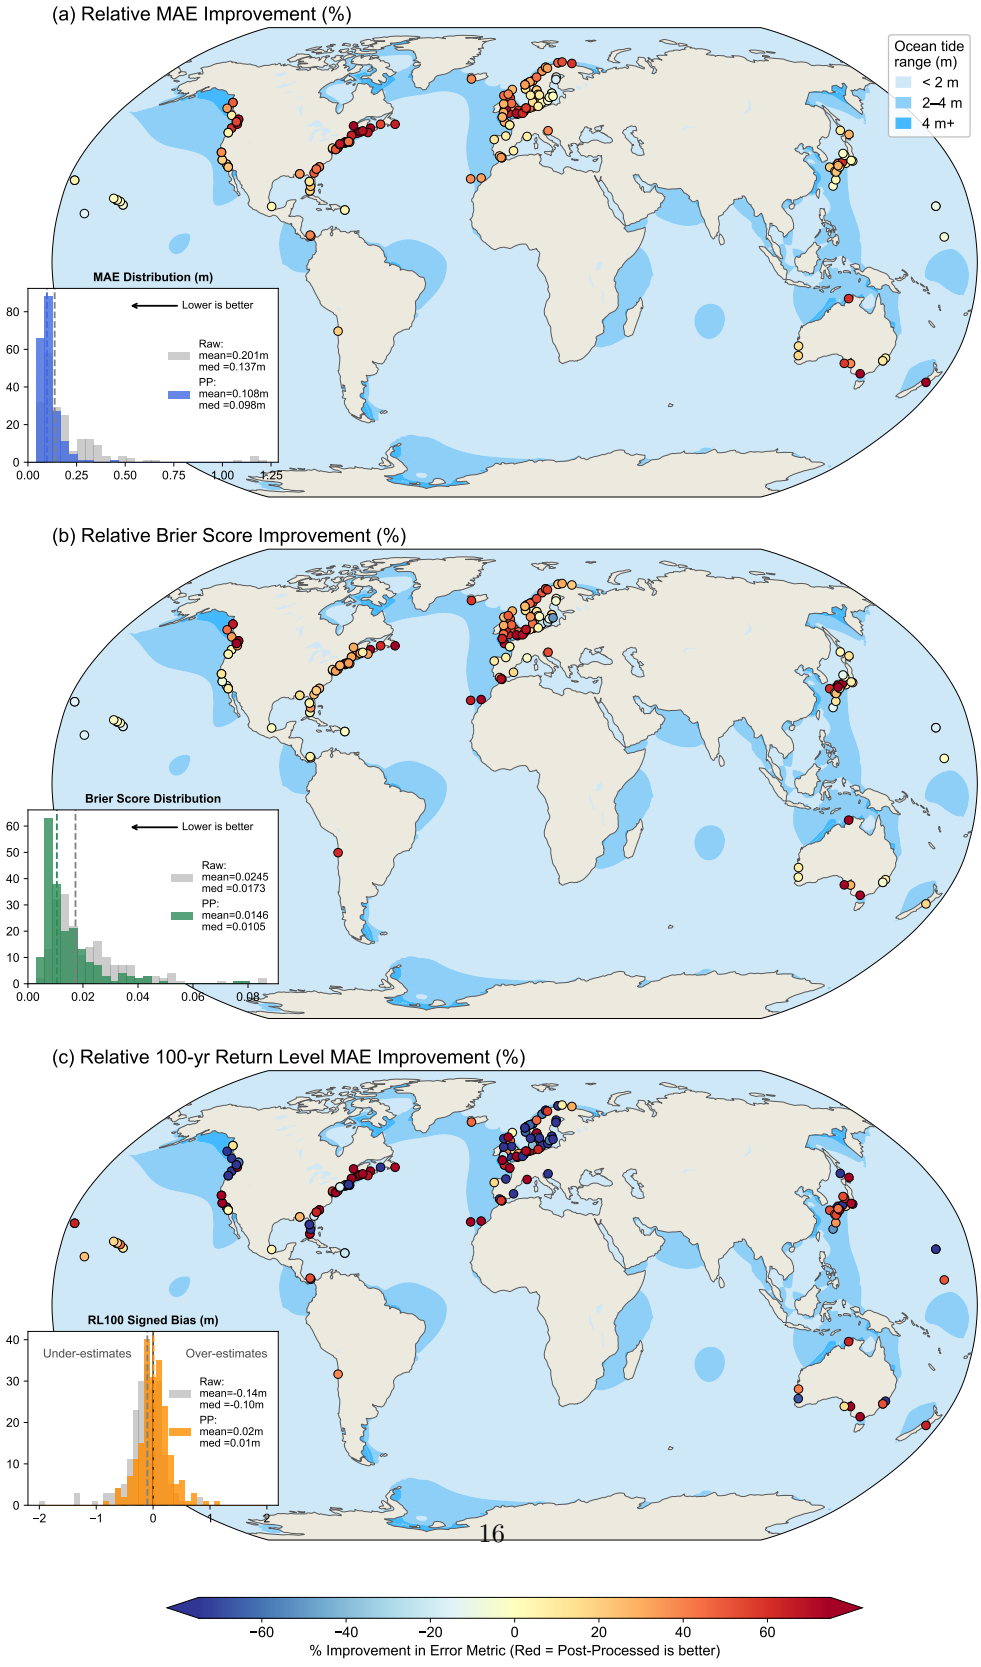

Fig. 3 Global bilinear postprocessing of GTSM sea-level reanalysis at GESLA4 tide gauges with 50 years of data (199 sites). Relative improvement in (a) MAE, (b) Brier score, and (c) 100-year return level accuracy under a leave-one-year-out protocol. Insets show the global distributions of each metric before and after postprocessing. Red indicates improvement; Blue indicates degradation. Ocean tidal range is shown in the background and computed from the FES2022 atlas.

network. Furthermore, the 20% reduction in 100-year return level MAE from 0.25 m to 0.2 m illustrates that a sizable amount of the uncertainty in extreme sea levels comes from high-frequency tidal processes. Using the GHS-POP 2025 population dataset [31], approximately 830 million people live within 50 km of a gauge. While this figure reflects proximity rather than direct flood exposure, it illustrates the potential societal reach of the approach. The concentration of gauges in the Northern Hemisphere (2118 of 2559) means that 648 million (78.1%) of this population lies in the north, with 182 million (21.9%) in the south. The geographic pattern of improvement is itself informative: the largest gains occur along dynamically complex coastlines where nonlinear tidal processes and unresolved bathymetry are prevalent, while degradation is concentrated at open-ocean and low-tidal-range sites where the model's structured residuals are small and there is little for the operator to learn. Residual limitations reflect the uneven global distribution of tide gauges [28]: coverage is sparse across the Global South, small island states, and the Arctic, where climate vulnerability is highest and the approach would be most valuable.

The time-invariant formulation enables exploitation of the full historical tide-gauge record and, as shown, the bilinear response associated with tidal processes remains robust for both mean and extreme estimation decades out of sample. The principal vulnerability of a time-invariant operator is nonstationarity: under climate change, shifts in mean sea level, storm climatology, and river discharge may alter the structured residuals the operator is trained to correct. This is partially mitigated by our approach: learning an operator that learns a mapping from total modeled sea-level to actual sea-level, rather than seeking to isolate individual processes, allows the response operator to "extrapolate" to unseen conditions or extremes, assuming those processes are included in the model. The trade-off, however, is that if these processes are not included, this extrapolation will fail. Periodic retraining on recent data is a natural mitigation, and the modest data requirement makes this practical.

Future work should explore the application of the same framework to both conventional and wide-swath altimetry data, particularly from the Surface Water Ocean Topography mission (SWOT) [32], to extend coverage to ungauged and unsurveyed regions. In this way, each altimetry observation location would serve as a virtual tide gauge and can be treated exactly as we do here. The errors that dominate in these regions, namely unresolved bathymetry and unmodeled coastal geometry, were shown in experiments S4 and S5 as recoverable through bilinear correction. A key advantage over tidal harmonic analysis in this context is that the bilinear operator infers a broadband spectral filter rather than decomposing into individual tidal lines. Discretizing into tidal lines creates severe aliasing when applied to temporally sparse altimetry data; by contrast, exploiting the predominantly correct frequency content of the dynamical model, the broadband filter theoretically reduces the observational data requirements needed to constrain the correction. Preliminary results on synthetic data spanning 3-years of SWOT and other altimetry sampling periods (Supplementary Figure S17) indicate the capacity of both approaches for such applications. Due to the sparsity of the altimetric observations, it is likely modification to the lag spacing to reduce parameter counts would reduce cases where performance worsens.

The novelty of the proposed approach lies in its implicit use of the physical information already present in the dynamical model. We show that the representation of several shallow-water nonlinear processes is achievable in this framework. However, this imposes two primary limitations: (i) if a physical process is absent from the model and not time-invariant, the assumptions of the approach are violated, and (ii) the gauge must faithfully represent localized dynamics. With regard to the former, this was observed in the Canadian SHOP model after the transition from tidal-dominated to fluvially dominated regimes. A simpler example of this would be to apply the approach to a hydrodynamic model output in a tidal river that does not include the river outflow. Here, the time-varying exogenous forcing leads to a loss of time-invariance between the model and data. Hence, to overcome this, future work should look to include such forcing within this framework. This is naturally achievable within several non-stationary tidal process models, including both RTide and NS_Tide [9, 33]. The second point about the need for gauge representivity is at the heart of data-driven modeling. Where errors are present in the gauge record, the inferred response weights may be biased. The approach thus imposes a requirement on the user to assess the representivity of their data and to ensure its quality.

8 Methods

8.1 Variational Bayesian weight estimation

We adopt a Bayesian linear modelling framework to infer the weights and associated hyperparameters of the postprocessing operator. A detailed derivation of variational inference can be found in [34], and a fuller exposition of the Bayesian linear model is provided in [15].

Given observations $Y = [y_0, y_1, \ldots, y_N]^\mathsf{T}$ and design matrix X, we seek to infer the posterior distribution over parameters and hyperparameters θ. This is obtained via Bayes' theorem,

$$p(\theta|Y) = \frac{p(\theta)\, p(Y|X,\theta)}{p(Y)}, \tag{8}$$

where $p(\theta)$ is the prior, $p(Y|X,\theta)$ the likelihood, and $p(Y)$ the marginal likelihood. The parameter set θ includes both model coefficients and hyperparameters governing their distributions.

Likelihood. We assume a Gaussian likelihood, corresponding to a least-squares error model,

$$p(Y|X,\theta) = p(Y|X,w,\beta) = \left(\frac{\beta}{2\pi}\right)^{N/2} \exp\left(-\frac{\beta}{2} E_Y(w)\right), \tag{9}$$

where β is the noise precision (inverse variance) and

$$E_Y(w) = \sum_{i=0}^{N} (y_i - w^\mathsf{T} x_i)^2. \tag{10}$$

The scalar precision parameter β controls the relative weighting of model–data misfit, allowing the model to adapt to varying observational noise levels.

Priors. We employ a hierarchical prior structure of the form

$$p(\theta) = p(w|\alpha)\, p(\alpha)\, p(\beta), \tag{11}$$

where $\alpha = \{\alpha_j\}$ are precision hyperparameters governing the scale of each model coefficient.

The weights are assigned independent zero-mean Gaussian priors,

$$p(w_j|\alpha_j) = \left(\frac{\alpha_j}{2\pi}\right)^{1/2} \exp\left(-\frac{\alpha_j}{2} w_j^2\right), \tag{12}$$

which allows coefficients to take either sign while penalizing large magnitudes. This formulation corresponds to Automatic Relevance Determination (ARD), whereby coefficients unsupported by the data are driven toward zero. This is an essential feature for handling the bilinear features which can have a high number of interaction terms.

The hyperparameters α_j are assigned Gamma priors,

$$p(\alpha_j) = \Gamma(\alpha_j; a_0, b_0), \tag{13}$$

with weakly informative hyperparameters $a_0 = 10^{-2}$ and $b_0 = 10^{-4}$. This choice promotes sparsity while maintaining flexibility in the inferred weight magnitudes.

The noise precision β is similarly assigned a Gamma prior,

$$p(\beta) = \Gamma(\beta; c_0, d_0), \tag{14}$$

with $c_0 = 10^{-2}$ and $d_0 = 10^{-4}$. This provides an adaptive estimate of observational uncertainty within the model.

Initialization. Model parameters are initialized using the maximum-likelihood solution,

$$w_{\mathrm{ML}} = (X^{\mathsf{T}} X)^{-1} X^{\mathsf{T}} Y, \tag{15}$$

which provides a consistent starting point for inference. The noise precision is initialized from the corresponding residual variance,

$$\beta^{-1} = \frac{1}{N} \sum_{i=1}^{N} (y_i - w_{\mathrm{ML}}^{\mathsf{T}} x_i)^2. \tag{16}$$

This initialization ensures that inference begins in a region of parameter space consistent with the data.

8.1.1 Variational inference

Exact Bayesian inference requires marginalization over the posterior, which is generally intractable for models of this form. Instead, we employ variational Bayes (VB), which approximates the true posterior with a tractable distribution.

We introduce an approximate posterior $q(\theta|Y)$ and adopt a mean-field factorization,

$$q(\theta|Y) = q(w|Y)\, q(\alpha|Y)\, q(\beta|Y), \tag{17}$$

which assumes independence between latent variables and enables efficient analytical updates.

The variational objective is obtained by decomposing the log evidence as

$$\log p(Y) = F(q) + \mathrm{KL}(q(\theta|Y) \,\|\, p(\theta|Y)), \tag{18}$$

where $F(q)$ is the evidence lower bound (ELBO) and $\mathrm{KL}(\cdot)$ is the Kullback–Leibler divergence between the approximate and true posteriors. Maximizing the ELBO is therefore equivalent to minimizing the divergence between q and the true posterior.

Under the factorization in Equation 17, the ELBO can be optimized by iteratively updating each factor in turn. This yields a set of closed-form update equations (see [35]) that are computationally efficient and scale to large datasets.

The resulting VB estimator provides approximate posterior distributions over all parameters and hyperparameters, enabling both uncertainty quantification and automatic complexity control through the ARD prior.

8.2 One-dimensional shallow water model

We employ a one-dimensional, cross-section-averaged shallow water model on a staggered C-grid to isolate individual tide–surge–fluvial interaction mechanisms in a controlled setting. The domain $x \in [0, L]$ runs from an open tidal boundary at $x = 0$ to either a closed wall or a fluvial inflow at $x = L$.

The continuity and momentum equations are

$$b\,\frac{\partial \eta}{\partial t} + \frac{\partial}{\partial x}\big(b\,\tilde{H}\,u\big) = 0, \tag{19}$$

$$\frac{\partial u}{\partial t} + g\,\frac{\partial \eta}{\partial x} = \frac{\tau_w}{\rho\,\hat{H}} - C_d(x)\,\frac{|u|\,u}{\hat{H}}, \tag{20}$$

where η is the free-surface elevation, u the depth-averaged velocity, $b(x)$ the channel width, g gravitational acceleration, ρ the water density, $\tau_w(t)$ the along-channel wind stress, and C_d the drag coefficient (described below). The total water depth is $H = h(x) + \eta$, with $h(x)$ the undisturbed bed depth. The depth variables $\tilde{H}$ and $\hat{H}$ in equations (19) and (20) are set according to the scenario under study (Table 1), enabling each nonlinear interaction pathway to be switched on or off independently.

Shelf configuration. In the default shelf configuration, the bed is flat ($h = h_0$), the width is uniform ($b = 1$, so that width cancels from the equations), and the landward boundary is a closed wall. This isolates the tide–surge interaction mechanisms (scenarios S1–S4) in the simplest possible geometry.

Estuary configuration. To investigate fluvial–tidal interactions and the influence of unresolved geometry, we additionally consider an estuarine domain in which the bed

depth tapers linearly from h_0 at the mouth to h_r at the head,

$$h(x) = h_0 + (h_r - h_0)\,\frac{x}{L}, \tag{21}$$

and the width converges exponentially,

$$b(x) = b_0\,\exp\!\big(-x/L_b\big), \qquad L_b = -\frac{L}{\ln(b_L/b_0)}, \tag{22}$$

where b_0 is the mouth width, b_L the head width, and L_b the e-folding convergence length. The landward boundary is replaced by a prescribed river discharge Q_r (see Supplementary Material). The drag coefficient is enhanced in shallow water following a Manning–Strickler-type scaling,

$$C_d(x) = C_{d,0}\,\max\!\Big(1,\ \frac{h_0}{h(x)}\Big), \tag{23}$$

which naturally increases friction in the shallow estuarine head where overtide generation is strongest. On the flat shelf ($h = h_0$), this reduces to a uniform $C_d = C_{d,0}$.

Two configurations are supported. In the *shelf* configuration, $h(x) = h_0$ is uniform and $b = 1$ (i.e. width plays no role). In the *estuary* configuration, the bed depth tapers linearly from h_0 at the mouth to h_r at the head,

$$h(x) = h_0 + (h_r - h_0)\,\frac{x}{L}, \tag{24}$$

and the width converges exponentially,

$$b(x) = b_0\,\exp\!\big(-x/L_b\big), \tag{25}$$

where b_0 is the mouth width and $L_b = -L/\ln(b_L/b_0)$ is the e-folding convergence length chosen so that $b(L) = b_L$.

Boundary conditions

At the seaward boundary, a Flather radiation condition is imposed,

$$u\big|_{x=0} = \frac{1}{2}\left(\frac{2\,c_0}{H_0}\,\eta_{\mathrm{bc}} + u\big|_{x=\Delta x} - \frac{c_0}{H_0}\,\eta\big|_{x=0}\right), \tag{26}$$

where $c_0 = \sqrt{gH_0}$ and the prescribed tidal elevation is

$$\eta_{\mathrm{bc}}(t) = r(t)[A\cos(\omega_{\mathrm{M}_2}\,t) + 0.35\,A\cos(\omega_{\mathrm{S}_2}\,t)]\,, \tag{27}$$

with a cosine ramp $r(t)$ over two M_2 periods to suppress start-up transients. The M_2 and S_2 periods are 12.42 h and 12.00 h, respectively; their superposition produces a $\sim$14.8-day spring–neap cycle.

At the landward boundary, either a solid wall ($u|_{x=L} = 0$) or a prescribed river discharge is applied:

$$u\big|_{x=L} = -\frac{Q_r(t)}{b(L)\,H(L,t)}, \tag{28}$$

with $Q_r > 0$ directed seaward.

A quadratic sponge layer occupying the first 15% of the domain near $x = 0$ relaxes η toward η_{bc} and applies Rayleigh damping to u, absorbing nonlinearly generated overtides (M_4, M_6) that would otherwise partially reflect from the linearised radiation condition.

Wind forcing Storm events are parameterised as time-limited wind stress pulses with a $\sin^2$ temporal envelope:

$$\tau_w(t) = \tau_0 \,\sin^2\!\left(\pi\,\frac{t - t_s}{T_s}\right)\cos\theta, \qquad t_s \le t \le t_s + T_s, \tag{29}$$

where τ_0 is the peak stress, T_s the storm duration, and θ the angle between the wind and the along-channel direction ($\tau_w = 0$ for cross-channel or offshore winds). Multiple storms are placed randomly with enforced minimum gaps.

Numerical method

The equations are discretised on a uniform staggered grid (N_x cells, spacing $\Delta x = L/N_x$) with η at cell centres and u at faces. Time integration uses a forward Euler step with *semi-implicit* treatment of the quadratic friction term: linearising the drag at the new time level gives

$$u_i^{n+1} = \frac{u_i^n + \Delta t\left(-g\,\dfrac{\partial \eta}{\partial x}\Big|_i + \dfrac{\tau_w}{\rho\,\hat{H}_i}\right)}{1 + \Delta t\, C_d\,|u_i^n|\,/\,\hat{H}_i + \gamma_i}, \tag{30}$$

where γ_i is the sponge damping coefficient (zero outside the sponge layer). This is unconditionally stable with respect to the friction term regardless of Δt. In the momentum equation, $1/\hat{H}$ is regularised as $\hat{H}/(\hat{H}^2+\delta^2)$ with $\delta = H_{\min}$ to prevent singularities near drying. A flux limiter in the continuity equation prevents H from falling below $H_{\min}$.

Default parameters Unless stated otherwise, the following values are used: $L = 100\,\mathrm{km}$ (shelf) or $60\,\mathrm{km}$ (estuary), $N_x = 400$, $\Delta t = 7.5\,\mathrm{s}$, $h_0 = 12\,\mathrm{m}$, $h_r = 3\,\mathrm{m}$, $b_0 = 5000\,\mathrm{m}$, $b_L = 200\,\mathrm{m}$, $A = 4\,\mathrm{m}$, $\rho = 1025\,\mathrm{kg\,m^{-3}}$, $C_d = 0.0025$, $H_{\min} = 0.5\,\mathrm{m}$. Storm parameters are drawn from $\tau_0 \sim \mathcal{N}(1.0,\, 0.3^2)\,\mathrm{Pa}$ (clipped at $0.05\,\mathrm{Pa}$), $T_s \sim \mathcal{N}(18,\, 4.5^2)\,\mathrm{h}$ (clipped at $1\,\mathrm{h}$), and $\theta \sim \mathcal{N}(0,\, 45)$.

8.3 Models and in situ data

UK NEMO Surge Models

Section 4 makes use of two versions of the UK Met Office Atlantic Margin Model AMM15 and AMM75. Both models use a three-dimensional NEMO configuration [21] of the Northwest European shelf seas [22, 23], using updated codebase [36] and

ERA5 reanalysis forcing [37, 38]. AMM15 has a fixed grid resolution of 1.5 km, and AMM75 sits on a downsampled version of the same grid with resolution 7.5 km (5x coarser). AMM15 resolves coastal dynamics at eddy-permitting scales and includes 51 terrain-following vertical levels, providing a physically consistent representation of tide–surge interactions. The same physical processes are represented in AMM75, albeit at lower resolution. Modeled total water levels are combined with concurrent hourly observational data taken from 43 class-A tide gauges around the UK Tide Gauge Network from 1993-2015.

DCSM7

Water level forecasts are of direct relevance to the management and maintenance of storm surge barriers in the Netherlands, and to the timely and correct closure of these barriers during storm surge events. Water level forecasts are therefore produced multiple times (4×) per day at the Hydro Meteo Centre. An important basis for these forecasts is the combination of the meteorological model ECMWF IFS and the hydrodynamic model DCSM7-100m.

The DCSM7-100m model covers the continental shelf between 15° West and 13° East longitude, and from 43° N to 64° North latitude. It is a so-called Flexible Mesh model with a grid resolution of 4 nautical miles in deep water, far from the coast in the western and north-western part of the model domain, with a gradual refinement of the grid to 0.25 nautical miles and 200 metres in the southern North Sea, and to 100 metres at the Dutch coast.

The output of the DCSM7-100m model is subsequently used as input for even higher-resolution water level and wave models in the vicinity of the Dutch storm surge barriers, for example, in the computation of hydraulic loads on the 62 gates of the Eastern Scheldt Barrier (Oosterscheldekering).

Verification studies have shown that forecasts from various models contain systematic model errors. This is also observed in the DCSM7-100m model. On the one hand, bias corrections are applied; however, spectral analysis of the error signals also reveals peaks in the amplitude density spectrum, which are in all likelihood attributable to highly localised geometric influences that cannot be, or cannot sufficiently be, reproduced in the numerical model computations. An example of this is shown in Figure S13 in the Supplementary Information, in which distinct peaks are clearly visible in the amplitude density spectrum of the error signal (model minus measurement) at the Antwerp gauging station.

In this study three gauges are utilized: Roompot Buiten, and Hoek van Holland in the Netherlands, and Sheerness, UK. These sites are chosen as they are used for the operational of three different storm surge barriers and are thus operationally relevant.

Stevens Flood Advisory System

The Stevens Flood Advisory System (SFAS) operates as an automated, real-time ensemble forecasting framework designed to predict water levels along the U.S. Mid-Atlantic and Northeast coasts. The system is executed four times daily, and each cycle produces water level estimates spanning approximately five and a half days. These outputs are partitioned into a one-day hindcast segment, used for model adjustment and

evaluation, followed by a four-and-a-half-day forecast horizon. The hydrodynamic and wave simulations within SFAS are driven by the Stevens Estuarine and Coastal Ocean Model [39–41] (sECOM), which is a free-surface hydrostatic numerical model that uses vertical coordinates (sigma) that follow the terrain with an orthogonal curvilinear Arakawa C-grid. Freshwater inflows into the coastal ocean and estuarine systems are incorporated using the output from the HEC-HMS developed by the U.S. Army Corps of Engineers, along with discharge forecasts provided by the NOAA Advanced Hydrologic Prediction Service [41]. To further enhance predictive accuracy, SFAS applies a tide bias correction based on multi-year historical datasets. In this procedure, harmonic analysis is applied to both modeled and observed water level time series at each station. The discrepancy between these reconstructed tidal signals is used to derive a correction term, which is subsequently applied to future forecasts [39, 42]. The uncertainty in atmospheric forcing is addressed through an ensemble approach, where sECOM is executed using 96 distinct meteorological forcing scenarios. The ensemble members are then combined using a weighted averaging scheme based on the root mean square error, assigning greater influence to simulations with lower errors to estimate total water levels during the hindcast period. Due to the limited availability of data (one year)[26], the evaluation presented in Section 5 uses the first six months of data for training and the remaining six months for tests.

Canadian SHOP

The Canadian Operational Hydrodynamic Prediction System (OHPS/SHOP) is a two-dimensional (2D) hydrodynamic modeling system developed for the St. Lawrence River and its fluvial estuary (https://eccc-msc.github.io/open-data/msc-data/nwp_ohps/readme_ohps_en/). The computational domain spans approximately 400 km, extending from the Montreal archipelago downstream through the fluvial estuary and into the estuarine transition zone. The domain can be schematically divided into four sections [43, 44]:

1. Montreal archipelago, where water level variability is primarily controlled by outflow from Lake Ontario and discharge from the Ottawa River.
2. Varennes to Port-Saint-François, where water levels are mainly influenced by seasonal hydrological cycles and long-term discharge variability. This region includes Lac Saint-Pierre, as well as major tributaries such as the Richelieu, Yamaska, Saint-François, and Saint-Maurice Rivers.
3. Trois-Rivières to Deschaillons, where water level variability reflects the combined influence of tides and river discharge variability across daily to seasonal timescales. Several tributaries, primarily along the north shore, contribute to the flow.
4. Portneuf to Saint-Joseph-de-la-Rive, where variability is dominated by tides and storm surges and where waters transition from fresh to brackish conditions past the Île d'Orléans.

The model is based on an unstructured grid with spatial resolution ranging from a few meters to several hundred meters (mean resolution 190 m). It provides near-real-time analyses (hindcasts) and 48-hour forecasts of water levels and depth-averaged

velocities, updated four times daily (00, 06, 12, and 18 UTC) at the Canadian Meteorological Centre.

SHOP is forced at its upstream boundaries and tributaries by hourly river discharge from the Water Cycle Prediction System (WCPS), using both observed and forecasted inputs. At the downstream (coastal) boundary, it is forced by observed or predicted tides and storm surges from the Regional Ensemble Storm Prediction System (RESPS). Atmospheric forcing—including wind, air temperature (for ice processes), precipitation, and surface evaporation—is provided by the High-Resolution Deterministic Prediction System (HRDPS). Seasonal processes are represented through prescribed ice cover maps (three classes of extent and thickness) during winter as well as macrophyte coverage maps that modulate bottom friction during the growth season [45].

This system supports operational applications in emergency management (e.g., oil spill response and flood forecasting), environmental prediction, and navigation. Operational monitoring is conducted using a network of 15 water level gauges distributed along the fluvial–estuarine continuum. In this study, observations from these stations are obtained from Fisheries and Oceans Canada (DFO) (https://tides.gc.ca/en/stations) at a 3-minute temporal resolution, along with concurrent model outputs at collocated grid points, for the period 2022-01-01 to 2024-12-31. All data are referenced to the CGVD28 vertical datum.

Although the hydrodynamic model has been developed and extensively validated against water level and velocity observations [46] using higher-resolution configurations [47–49], the operational implementation employs a coarser grid to ensure computational efficiency. This introduces representation errors in regions with unresolved topographic features or sharp bathymetric gradients. Additional uncertainties arise from boundary forcing, including errors in upstream discharge and downstream storm surge predictions, which propagate throughout the domain.

To mitigate these effects, a minimal postprocessing step consisting of offline bias correction at tide gauge locations is currently applied prior to dissemination by DFO. However, this approach remains limited in its ability to correct temporally correlated errors over the full forecast horizon as well as spatially distributed errors, motivating the development of more advanced postprocessing techniques.

MIKE21 Tidal Currents

The Pentland Firth is characterized by some of the fastest tidal currents in the world and is thus of high-interest for tidal stream energy site development []. Presently, the MeyGen site, located in the Inner Sound of the Pentland Firth, is the largest tidal stream energy installation in the world []. We utilize data from the operational depth-average MIKE21 hydrodynamic model used at the site [50]. This model is employed for both site assessment flow characterization, and operational forecasting for energy yield. Due to the proprietary nature of the model, exact implementation details are not available, but it is expected that its usage in operational practice reflects an adherence to best practices and calibration with localized data. Concurrent current velocity data for the year 2020 were recorded via a seabed-mounted ADCP in the Inner Sound, processed using 10-minute averaging and standard outlier rejection protocols. Since

only one year of data is available, testing in Section 5 uses the first 6 months of data for training and the remaining 6 months for tests.

GTSM ERA5-E

We use the GTSMv3.0 global tide-and-surge reanalysis configuration developed and made available by [30]. We here provide a condensed summary but interested readers are referred to the original text for comprehensive implementation details in benchmarking.

The model is implemented in Delft3D Flexible Mesh [51] as a depth-averaged hydrodynamic model on an unstructured grid with spatial resolution reaching about 2.5 km along most global coasts and 1.25 km along the European coast. The model is forced with hourly ERA5 10 m winds and mean sea-level pressure, includes a spatially varying annual mean sea-level field, has no open boundaries, and represents tides through 60 tidal constituents while storm surges arise from wind stress and inverse-barometer forcing using a Charnock drag formulation. Bathymetry is assembled from GEBCO globally, EMODnet in European waters, and Bedmap2 beneath Antarctic ice shelves. Here, we only make use of the model estimated still water levels.

We combine the model outputs with GESLA4 tide gauges [28]. GESLA4 is the latest generation of the Global Extreme Sea Level database and provides an extensive inventory of global tide gauges. Similar to the analysis in [30], we first identify gauges greater than 50 years data coverage, with 75% record completion, between 1950 and 2025. We collocate these gauges with GTSM grid-points and select gauges within 5 km. Prior to analysis, we identified multiple gauges with datum shifts, which led to inaccurate estimates of model performance (both for GTSM and the post-processed variants). The datum shift identification procedure is described in the Supplementary Information. After removing these gauges and accounting for duplicate gauges, our analysis is carried out on 199 GESLA4 gauges.

8.4 Metrics

Mean Absolute Error

The mean absolute error (MAE) is defined as

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N} |\hat{y}_i - y_i| \,, \tag{31}$$

where $\hat{y}_i$ denotes the model prediction, y_i the corresponding observation, and N the number of samples. Lower values indicate better agreement between model and observations.

Brier score at the 99th percentile.

To assess extreme-value skill, we define a binary exceedance event at the 99th percentile of the training observations,

$$u_{99} = Q_{0.99}\big(\zeta \,\big|\, \mathrm{train}\big), \tag{32}$$

which fixes the threshold independently of the test data.

The Brier score is then given by

$$\mathrm{BS}_{99} = \frac{1}{N}\sum_{i=1}^{N}\left[\mathbf{1}\Big(\hat{\zeta}(t_i) > u_{99}\Big) - \mathbf{1}(\zeta(t_i) > u_{99})\right]^2, \tag{33}$$

where $\hat{\zeta}(t_i)$ denotes the model prediction, $\zeta(t_i)$ the corresponding observation, and $\mathbf{1}(\cdot)$ is the indicator function. Here, the indicator function simply maps the continuous values of ζ to a binary outcome variable such that

$$\mathbf{1}(A) = \begin{cases} 1, & \text{if } A \text{ is true,} \\ 0, & \text{otherwise.} \end{cases}$$

Because we only consider deterministic predictions (neglecting uncertainty inferred by the variational Bayes), Equation (33) reduces to the misclassification rate for threshold exceedance. Lower values indicate better agreement between predicted and observed extreme events. It is inappropriate to compare raw Brier scores across different data sets (e.g., comparing GTSM results with UK AMM- results). However, if two models are evaluated on the same data, Brier scores can be directly compared.

Percentage improvement.

To quantify the relative change in performance after postprocessing, we define the percentage improvement with respect to the baseline raw model output M_{raw}, as

$$\mathrm{PI}(\%) = 100 \times \frac{M_{\mathrm{raw}} - M_{\mathrm{post}}}{M_{\mathrm{raw}}}. \tag{34}$$

Here, M_{post} denotes the metric computed from the postprocessed predictions. Because both MAE and Brier score are error metrics for which lower values indicate better performance, $\mathrm{PI} > 0$ indicates an improvement, while $\mathrm{PI} < 0$ indicates a degradation in performance.

8.5 Return-level estimation

We follow a similar methodology for estimating return levels as used in [30]. However, rather than demean and detrend our data using the entire data series, which was originally done to remove the impacts of fixed model biases and sea-level rise, we instead use the observational data to derive these offsets.

For each site, we fit a Generalised Pareto Distribution (GPD) to peaks over a 99th-percentile threshold u [2]. Exceedances within 72 h of each other are declustered by retaining only the local maximum, and the GPD is fitted to the cluster excesses $(Y_i - u)$ with location fixed at zero using `scipy.stats.genpareto.fit` [52]. Sites yielding fewer than 10 cluster peaks are discarded. Writing λ for the mean annual rate

of cluster exceedances over the record length, the T-year return level is

$$\mathrm{RL}(T) \;=\; u \;+\; \frac{\hat{\sigma}}{\hat{\xi}}\Big[(T\lambda)^{\hat{\xi}} - 1\Big], \tag{35}$$

evaluated at $T \in \{10, 50, 100\}$ yr. The GPD is fitted independently to three pooled *out-of-sample* test-fold series at each site: (i) the observations ζ, (ii) the post-processed predictions $\hat{\zeta}_{\mathrm{PP}}$, and (iii) a bias-corrected GTSM baseline $\hat{\zeta}_{\mathrm{base}} = x_{\mathrm{GTSM}} - \Delta$, where

$$\Delta \;=\; \overline{x_{\mathrm{GTSM}}}\big|_{\mathrm{train}} - \overline{\zeta}\big|_{\mathrm{train}} \tag{36}$$

is the mean offset computed on the training period only. This offset is essential for fair comparison: the regression implicitly learns an intercept from the training period, so granting the raw baseline the same training-period mean calibration isolates the contribution of lagged and nonlinear structure from simple mean-level correction. No test-period information is used to set Δ.

The reported return-level bias is the absolute difference

$$\mathrm{Bias}^{\mathrm{RL}}(T) \;=\; \mathrm{RL}_{\mathrm{model}}(T) - \mathrm{RL}_{\mathrm{obs}}(T), \tag{37}$$

computed separately for the post-processed and bias-corrected baseline series, with $\mathrm{MAPE}^{\mathrm{RL}}(T) = 100\,|\mathrm{Bias}^{\mathrm{RL}}(T)|/|\mathrm{RL}_{\mathrm{obs}}(T)|$. The return-level analysis in Equation (35) does *not* apply this test-period correction — it uses only the training-period Δ of Equation (36) — so that the GPD is fit to series that could in principle be produced operationally.

Supplementary information. Supplementary information including additional results and discussion on the theoretical capacity of the linear and bilinear operators (Figures S1-6), operational results on the Canadian SHOP model (Figures S7-8), further evaluation of the global performance on GTSM and GESLA4 (Figures S9-10) and comparisons of variational Bayes with other estimators (Figures S11-12) is provided in a separate file.

Acknowledgements. KS and TM acknowledge the I-STORM forecasting network for facilitating knowledge sharing and collaboration. JP acknowledges the use of ARCHER2 UK National Supercomputing Service (http://www.archer2.ac.uk), Joanne Williams (who performed the surge QC on the tide gauge data), and the UK Joint Marine Modelling Programme (which facilitates the operational model development). TM also acknowledges Xinyu Liu for fruitful discussions.

- **Funding** TM would like to acknowledge funding from the Eric and Wendy Schmidt Foundation, as well as New College at the University of Oxford. This work was also supported by the Natural Environment Research Council (NERC) [grant number NE/W004976/1] as part of the Agile Initiative at the Oxford Martin School. JP was funded under the NERC AtlantiS project (NE/Y005589/1).
- **Conflict of interest** The authors have no conflicts of interest to declare.

- **Data Availability** Data are available from the authors upon request. The GTSM extended reanalysis [30] can be accessed freely through [53]. Data in this manuscript was accessed via the Copernicus Climate Data Service: [54]. GESLA4 data can be accessed through https://gesla787883612.wordpress.com/. Ensemble data from the Steven's Institute forecast can be downloaded freely https://data.mendeley.com/datasets/7mx8npghk9/1 [26]. Tide gauge data from The Barrier can be accessed freely via the NOAA Tides and Currents website: https://tidesandcurrents.noaa.gov/.
  Due to sensitivity issues with operational forecasts, UK AMM15 and AMM75 surge data, DCSM7, and Canadian SHOP models are available upon request. 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 an academic partnership for the current study.
- **Code availability** Code is available from the authors upon request. After publication, code will be placed into a permanent repository. Weights associated with the 2559 GESLA4 sites identified with 1-year of data will also be made available in the final repository.
- **Author contribution** TM conceptualized the study and methodology, managed product administration, and conducted the formal investigation including software development. J.P., P.M., and S.I. contributed to the investigation design and formal analysis interpretation. TM and TAAA contributed to funding acquisition, and TAAA provided formal supervision. All authors contributed various data as well as to writing (original draft, review, and editing).

References

[1] Haigh, I. D. & Nicholls, R. J. Coastal flooding. *MCCIP Science Review 2017* 98–104 (2017).

[2] Wahl, T. *et al.* Understanding extreme sea levels for broad-scale coastal impact and adaptation analysis. *Nature communications* **8**, 16075 (2017).

[3] Green, J. *et al.* A comprehensive review of compound flooding literature with a focus on coastal and estuarine regions. *Natural Hazards and Earth System Sciences* **25**, 747–816 (2025).

[4] Holt, J. T., Allen, J. I., Proctor, R. & Gilbert, F. Error quantification of a high-resolution coupled hydrodynamic–ecosystem coastal–ocean model: part 1 model overview and assessment of the hydrodynamics. *Journal of Marine Systems* **57**, 167–188 (2005).

[5] Mejia-Olivares, C. J., Haigh, I. D., Lewis, M. J. & Neill, S. P. Sensitivity assessment of bathymetry and choice of tidal constituents on tidal-stream energy resource characterisation in the gulf of california, mexico. *Applied Ocean Research* **102**, 102281 (2020).

[6] Bosboom, J. & Stive, M. J. *Coastal dynamics* (TU Delft OPEN Publishing, 2023).

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

[8] Cartwright, D. E. 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**, 1–55 (1968).

[9] Monahan, T., Tang, T., Roberts, S. & Adcock, T. A. A. Rtide: Automating the tidal response method. *Journal of Geophysical Research: Machine Learning and Computation* **2**, e2024JH000525 (2025).

[10] Matte, P. & Innocenti, S. Filling gaps and extending inland tide gauge records to support accurate long-term and extreme water level analyses. *Continental Shelf Research* 105668 (2026).

[11] Giaremis, S. *et al.* Storm surge modeling in the ai era: Using lstm-based machine learning for enhancing forecasting accuracy. *Coastal Engineering* **191**, 104532 (2024).

[12] Tedesco, P. *et al.* Bias correction of operational storm surge forecasts using neural networks. *Ocean Modelling* **188**, 102334 (2024).

[13] Nader, N. *et al.* Storm surge modeling, bias correction, graph neural networks, graph convolution networks. *arXiv preprint arXiv:2604.20688* (2026).

[14] Zetler, B. D. & Munk, W. H. The optimum wiggliness of tidal admittances (1975).

[15] Roberts, S., McQuillan, A., Reece, S. & Aigrain, S. Astrophysically robust systematics removal using variational inference: application to the first month of Kepler data. *Monthly Notices of the Royal Astronomical Society* **435**, 3639–3653 (2013).

[16] Monahan, T., Tang, T., Roberts, S. & Adcock, T. A. A. Tidal corrections from and for swot using a spatially coherent variational bayesian harmonic analysis. *Journal of Geophysical Research: Oceans* **130**, e2024JC021533 (2025).

[17] Niyazi, Y. *et al.* Status of global seafloor mapping effort and priority areas for future mapping. *Frontiers in Marine Science* **12**, 1543885 (2025).

[18] Jay, D. A. Green's law revisited: Tidal long-wave propagation in channels with strong topography. *Journal of Geophysical Research: Oceans* **96**, 20585–20598 (1991).

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

[20] Godin, G. The propagation of tides up rivers with special considerations on the upper saint lawrence river. *Estuarine, Coastal and Shelf Science* **48**, 307–324 (1999).

[21] Madec, G. *et al.* Nemo ocean engine (2015).

[22] Graham, J. A. *et al.* Amm15: a new high-resolution nemo configuration for operational simulation of the european north-west shelf. *Geoscientific Model Development* **11**, 681–696 (2018).

[23] Tonani, M. *et al.* The impact of a new high-resolution ocean model on the met office north-west european shelf forecasting system. *Ocean Science* **15**, 1133–1158 (2019).

[24] Uncles, R. Physical properties and processes in the bristol channel and severn estuary. *Marine Pollution Bulletin* **61**, 5–20 (2010).

[25] CNES. Fes2022 (finite element solution) tidal model (version 2024). Data set (2024). URL https://doi.org/10.24400/527896/A01-2024.004. The FES2022 Tide product was funded by CNES, produced by LEGOS, NOVELTIS and CLS and made freely available by AVISO.

[26] Chen, Z. The ensembles forecast data for the stevens flood advisory system (2026). URL https://doi.org/10.17632/7mx8npghk9.1.

[27] Stammer, D. *et al.* Accuracy assessment of global barotropic ocean tide models. *Reviews of geophysics* **52**, 243–282 (2014).

[28] Haigh, I. D. *et al.* Gesla version 4: Further enhancements to the global high-frequency sea-level dataset (2026). In preparation.

[29] Hart-Davis, M. *et al.* Tides in complex coastal regions: Early case studies from wide-swath swot measurements. *Geophysical Research Letters* **51**, e2024GL109983 (2024).

[30] Aleksandrova, N., Veenstra, J. & Muis, S. Global dataset of storm surges and extreme sea levels for 1950–2024 based on the era5 climate reanalysis. *Earth System Science Data Discussions* **2025**, 1–22 (2025).

[31] Schiavina, M., Freire, S., Carioli, A. & MacManus, K. GHS-POP R2023A - GHS population grid multitemporal (1975-2030) (2023). URL http://data.europa.eu/89h/2ff68a52-5b5b-4a22-8f40-c41da8332cfe. PID: http://data.europa.eu/89h/2ff68a52-5b5b-4a22-8f40-c41da8332cfe.

[32] Morrow, R. *et al.* Global observations of fine-scale ocean surface topography with the surface water and ocean topography (SWOT) mission. *Frontiers in Marine Science* **6**, 232 (2019).

[33] Matte, P., Jay, D. A. & Zaron, E. D. Adaptation of classical tidal harmonic analysis to nonstationary tides, with application to river tides. *Journal of Atmospheric and Oceanic Technology* **30**, 569–589 (2013).

[34] Fox, C. W. & Roberts, S. J. A tutorial on variational Bayesian inference. *Artificial intelligence review* **38**, 85–95 (2012).

[35] Penny, W. & Roberts, S. Bayesian multivariate autoregressive models with structured priors. *IEE Proceedings-Vision, Image and Signal Processing* **149**, 33–41 (2002).

[36] Patmore, R. *et al.* Jmmp-group/co_amm15: Co9_p2.0 (2023). URL https://doi.org/10.5281/zenodo.8171688.

[37] Hersbach, H. *et al.* The era5 global reanalysis. *Quarterly journal of the royal meteorological society* **146**, 1999–2049 (2020).

[38] Patmore, R. D. Noc-msm/co_amm15_chamfer: Initial release (2025). URL https://doi.org/10.5281/zenodo.14937563.

[39] Georgas, N. & Blumberg, A. F. in *Establishing confidence in marine forecast systems: The design and skill assessment of the new york harbor observation and prediction system, version 3 (nyhops v3)* 660–685 (2010).

[40] Orton, P., Georgas, N., Blumberg, A. & Pullen, J. Detailed modeling of recent severe storm tides in estuaries of the new york city region. *Journal of Geophysical Research: Oceans* **117** (2012).

[41] Georgas, N. *et al.* The stevens flood advisory system: Operational h3e flood forecasts for the greater new york/new jersey metropolitan region. *Flood Risk Management and Response* 194 (2016).

[42] Ayyad, M., Orton, P. M., El Safty, H., Chen, Z. & Hajj, M. R. Ensemble forecast for storm tide and resurgence from tropical cyclone isaias. *Weather and Climate Extremes* **38**, 100504 (2022).

[43] Matte, P., Secretan, Y. & Morin, J. 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**, 5724–5744 (2014).

[44] Innocenti, S. *et al.* Tidal, hydrological and meteorological contributions to high-water level events in the Saint Lawrence River Estuary: Local responses to regional drivers. *Weather and Climate Extremes* **51**, 100852 (2026).

Supporting information for Learning unresolved coastal dynamics in hydrodynamic models

Thomas Monahan[1*], Jeffrey Polton[2], Silvia Innocenti[3], Pascal Matte[3], Mahmoud Ayyad[4], Krijn Saman[5], Thomas A. A. Adcock[1]

[1*]Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ , United Kingdom.

*Corresponding author(s). E-mail(s): thomas.monahan@eng.ox.ac.uk; Contributing authors: jelt@noc.ac.uk; Silvia.Innocenti@ec.gc.ca; Pascal.Matte@ec.gc.ca; mayyad@stevens.edu; krijn.saman@rws.nl; thomas.adcock@eng.ox.ac.uk;

1 Overview

This document contains additional experiments and results that support the findings in the main text.

2 Lag configuration: theory and justification

The spacing $\Delta\tau$ and maximum lag L together fix (i) the frequency band over which the filter can express structure, (ii) the number of free parameters, and (iii) the volume of data required to estimate them. This section sets out why the choice $\Delta\tau = 1$ h, $L = 24$ h is a principled compromise for this problem. We do not explicitly discuss the long-period filter as the theory is identical.

2.1 Linear regression as an FIR filter

The linear component of the regression,

$$\hat{y}(t) \;=\; \sum_{k=0}^{n-1} w_k \, x(t - k\Delta\tau)\,, \tag{1}$$

is a finite impulse response (FIR) filter with $n = L/\Delta\tau + 1$ is the number of so-called "taps" at $\tau_k = k\Delta\tau$ and weights w_k. Its transfer function is

$$H(f) \;=\; \sum_{k=0}^{n-1} w_k \, e^{-i\,2\pi f \tau_k}, \tag{2}$$

so $|H(f)|$ is the admittance — the amplitude response of the filter at frequency f. Figure 1 plots $|H(f)|$ for uniform-weight filters across several $(\Delta\tau, L)$ choices and makes the consequences of the lag configuration explicit.

Upper frequency limit (Nyquist).

Because the taps are spaced by $\Delta\tau$, the filter is subject to the discrete-time Nyquist bound $f_{\mathrm{Nyq}} = 1/(2\Delta\tau)$ [1]. At $\Delta\tau = 1$ h this is 0.5 cycles-per-hour (cph), which comfortably covers the semidiurnal band ($M_2 \approx 0.0805$ cph, $S_2 \approx 0.0833$ cph) and the first four overtide harmonics M_4, M_6, M_8. Increasing the spacing to $\Delta\tau = 6$ h collapses f_{Nyq} to 0.083 cph, placing M_4 and everything above it into the aliased band (orange shading in Figure 1a). Because the nonlinear tide–surge interaction explicitly transfers energy into the overtides, any spacing that aliases this band is unfit for purpose regardless of how much data is available.

Lower frequency limit (window length).

A filter of maximum lag L cannot represent variability whose characteristic period exceeds L: signals slower than the window appear only as a flat DC offset. The lowest partially resolvable period is therefore $T_{\min} \approx L$. With $L = 24$ h we retain the diurnal band and all higher frequencies but deliberately exclude slower modulation (e.g. the MS_f fortnightly beat at $T \approx 354$ h, seasonal steric drift, or fluvial effects that are not explicitly included in the hydrodynamic model). These processes are not driven by instantaneous interactions and should explicitly be treated separately.

Frequency resolution.

The width of each spectral lobe in Figure 1(a) scales as $1/L$: longer windows give sharper discrimination and deeper nulls between frequency bands. Figure 2 summarises the two bounds together. The chosen configuration sits in the admissible region: f_{Nyq} well above M_8 and $T_{\min}$ at the diurnal period.

2.2 Parameter count and the cost of extending the window

The linear regression has $n = L/\Delta\tau + 1$ coefficients, but the second-order Volterra (bilinear) model adds a further $n(n+1)/2$ cross-terms, growing quadratically in n (Figure 3). Doubling L from 24 to 48 h at $\Delta\tau = 1$ h raises the linear count from 25 to 49, but the bilinear count from 325 to 1225. One year of hourly observations yields ~ 8760 samples, giving roughly 7 samples per parameter for the extended model — too few for stable identification, even under strong regularisation. The operational choice retains ~ 25 samples per parameter, which sits comfortably in the tractable regime when combined with the ARD prior that prunes inactive cross-terms. This

relationship motivates handling long-period processes separately as learning a joint filter which spans both frequency ranges becomes intractable.

2.3 Data requirements

Figure 4 shows the test-set mean absolute error versus training length for both models, with four levels of additive observation noise. The training length evaluations are carried out by training the models on increasingly long contiguous records. We do this by randomly selecting 40 periods of a given duration and then testing on the same out-of-sample 60 day test set. Reported values are the average across all folds. It can be seen that for both approaches increasing the amount of data yields improved out-of-sample performance. However, performance improvements diminish with increasing data length.

Figure 5 decomposes this convergence by frequency. Tidally forced lines (M_2, M_4) are recovered after ~ 20 days because the forcing spectrum concentrates energy at these frequencies. Intertidal frequencies need 60–120 days to stabilise: these admittance values are estimated from broadband residuals in which the signal-to-noise ratio is intrinsically low. The plateau in Figure 4 reflects this partitioning — the informative (tidally forced) part of the spectrum is learned early, and further data mostly refines weakly excited components that contribute little to total error. Where the non-tidal signal is comparatively stronger, we expect the model would identify these data earlier. This is an effect of the regularisation of the ARD priors which allows the model to effectively learn a data-driven filter which increases in complexity with data.

We further evaluate the stability of the learned bilinear interactions by studying the weighted quadratic transfer functions (QTF) in Figure 6. It can be seen that while increasing data reduces variability in the off-diagonal terms, the transfer functions remain qualitatively stable even using 30 days of data.

3 Investigation into tidal river dynamics

The St. Lawrence Estuary provides a unique test-bed with which to evaluate how the model performs before and after the transition from tide-surge dominated to fluvially dominated regimes. Figure 7 shows the residual autocorrelation function (ACF) computed between the modeled and observed water levels. A value of zero reflects that there is no correlation between the operational SHOP model and the gauge observation at that time-period. We observe that after Portneuf, there is a sharp transition in dynamics in which the ACF shows correlations extending beyond 30-d. This marks the transition between tide-surge dominated and fluvially reaches. This shift is reflected in the model performance which is shown in Figure 8. We observe a degradation in performance of the bilinear operator after this transition. This makes sense as the filter is explicitly designed to learn a broadband filter defined over the tidal spectrum. Given the dominant interactions are occuring on much larger time-scales, this filter has no utility here.

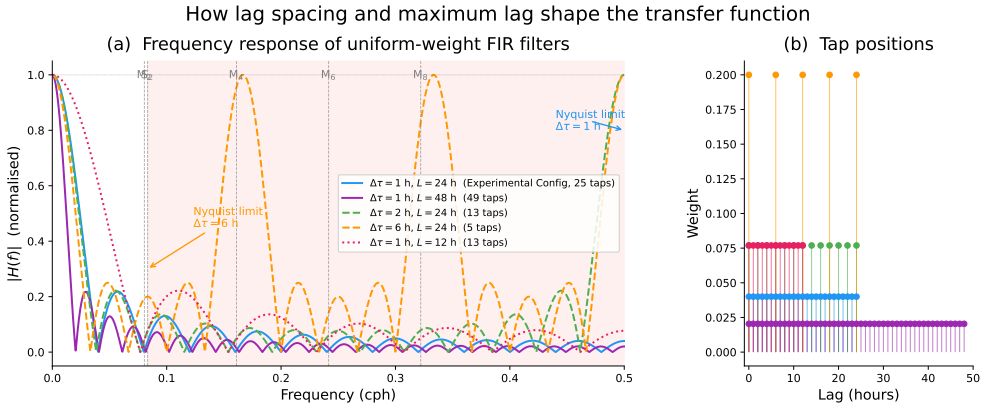

Supplementary Figure 1 Frequency response $|H(f)|$ of uniform-weight FIR filters for five lag configurations. **(a)** The spacing $\Delta\tau$ sets the Nyquist limit $f_{\mathrm{Nyq}} = 1/(2\Delta\tau)$; the coarse configuration $\Delta\tau = 6$ h (orange) folds the higher overtides into the aliased red-shaded region above $f_{\mathrm{Nyq}} = 0.083$ cph. The maximum lag L controls the width of each spectral lobe ($\propto 1/L$), so longer filters (purple, $L = 48$ h) produce sharper nulls between tidal harmonics. The operational configuration $(\Delta\tau, L) = (1, 24)$ h (blue) resolves M_2 through M_8 while keeping the parameter count modest. **(b)** Corresponding tap positions in the time domain.

4 Additional GTSM Results

4.1 Datum shift identification

To identify discrete vertical datum shifts in the residual time series, we implemented a scanning changepoint detection algorithm. Residuals from the uncorrected GTSM output were downsampled to daily means $\bar{R}_t$ to isolate persistent shifts from high-frequency noise. Note that, since this was carried out using the raw GTSM data, the removal of gauges does not artificially favor our proposed approach. For each site, we evaluated every potential split point k within the valid range $[L_{min}, N - L_{min}]$, where $L_{min} = 180$ days.

At each candidate cut, the series was partitioned into two segments (S_1, S_2) and evaluated using the Welch's t-statistic to account for unequal variances and sample sizes:

$$t(k) = \frac{\mu_2 - \mu_1}{\sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}} \tag{3}$$

The optimal changepoint k^* was defined as the index maximizing $|t(k)|$. Statistical significance was determined using the Welch–Satterthwaite approximation for effective degrees of freedom ν:

$$\nu \approx \frac{\left(\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}\right)^2}{\frac{(\sigma_1^2/n_1)^2}{n_1 - 1} + \frac{(\sigma_2^2/n_2)^2}{n_2 - 1}} \tag{4}$$

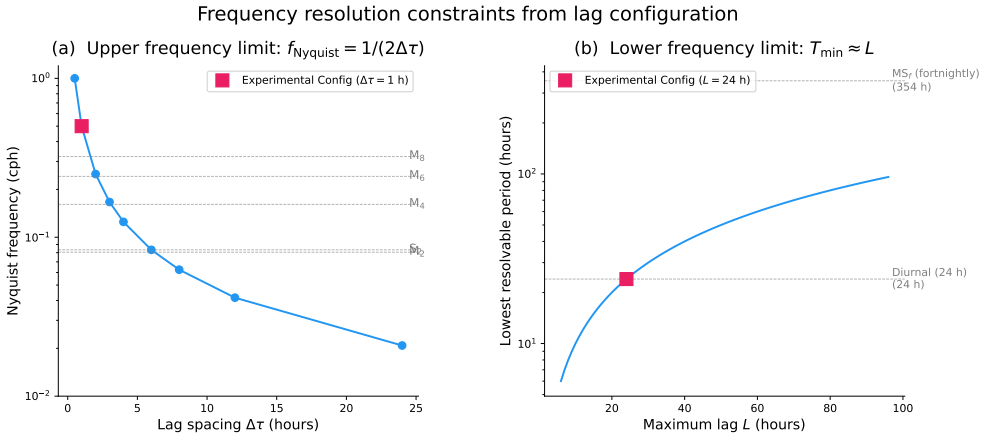

Supplementary Figure 2 Two complementary bounds on what the FIR filter can resolve. **(a)** Nyquist frequency $f_{\text{Nyq}} = 1/(2\Delta\tau)$ as a function of lag spacing; at $\Delta\tau = 1$ h (magenta square) f_{Nyq} lies well above all semidiurnal overtides of interest. **(b)** Lowest resolvable period $T_{\text{min}} \approx L$ as a function of maximum lag; at $L = 24$ h we resolve down to the diurnal band but deliberately exclude lower-frequency modulation such as the MS_f fortnightly beat. Together these two bounds define the spectral window within which the regression can express structure.

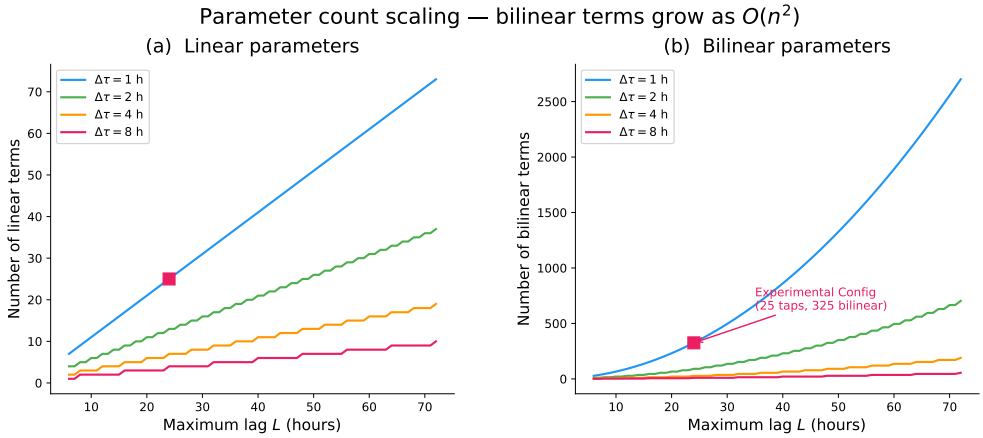

Supplementary Figure 3 Number of free parameters in the regression as a function of maximum lag L, for four spacings. **(a)** The linear parameter count $n = L/\Delta\tau + 1$ grows linearly. **(b)** The second-order Volterra (bilinear) parameter count $\binom{L^{(n)}+n}{n}$ grows quadratically even when enforcing symmetry, so doubling either L or $1/\Delta\tau$ quadruples the number of cross-terms to identify. The operational configuration (magenta marker, $\Delta\tau = 1$ h, $L = 24$ h) gives 25 linear and 325 bilinear terms (350 total), for which one year of hourly data provides roughly 25 samples per parameter — sufficient for stable identification under ARD regularisation.

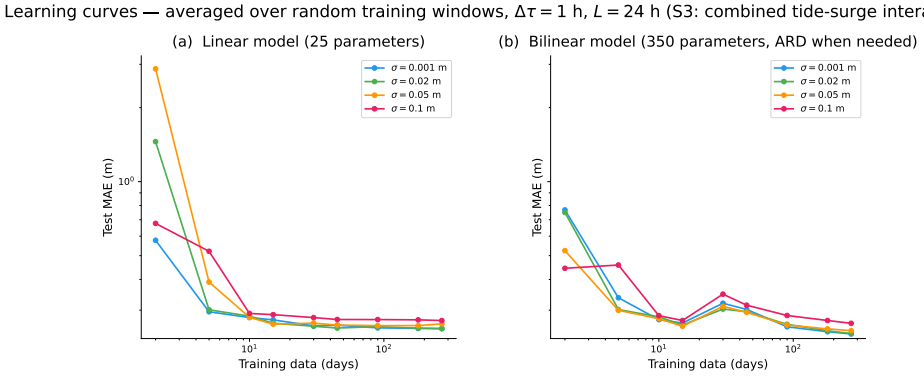

Supplementary Figure 4 Test-set mean absolute error against training-data length, using the synthetic shallow-water simulation (scenario S3, combined tide–surge interaction), evaluated under four levels of additive Gaussian observation noise. Dotted vertical line marks 60 days, the typical useable segment length of a single-year GESLA record. **(a)** The 25-parameter linear model converges smoothly and plateaus by $\sim$60 days, with negligible sensitivity to realistic noise levels ($\sigma \leq 0.05$ m). **(b)** The 350-parameter bilinear model with ARD regularisation is more data-hungry: MAE is non-monotonic in the 20–30 day range as ARD alternately admits and then rejects borderline cross-terms, but reaches lower error than the linear model once $\gtrsim$60 days are available.

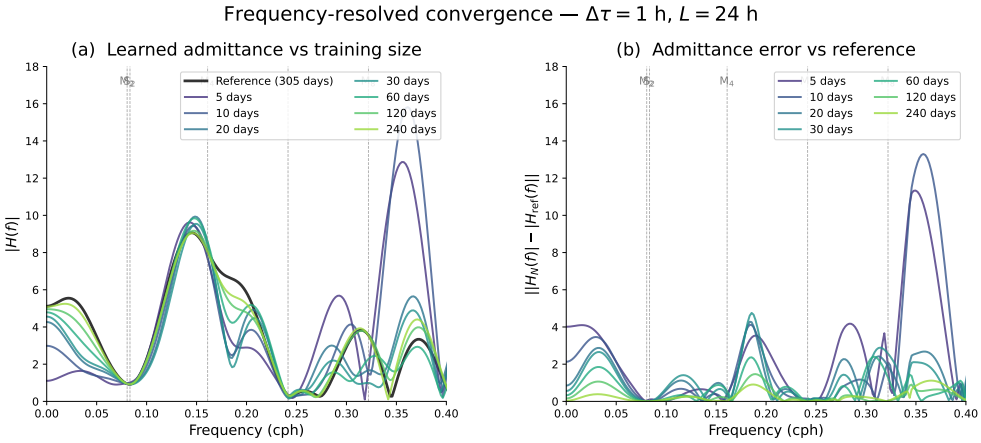

Supplementary Figure 5 Frequency-resolved convergence of the linear-model admittance towards its long-data limit. **(a)** $|H(f)|$ for training lengths of 5–240 days (viridis colour scale) overlaid on a 305-day reference fit (black). The dominant M_2 admittance peak near 0.08 cph and the M_4 overtide are essentially locked by 20 days, because the forcing spectrum concentrates energy at these lines. **(b)** Absolute error $\big||H_N(f)| - |H_{\mathrm{ref}}(f)|\big|$ as a function of training length. The residual is largest at intertidal frequencies and near the Nyquist edge (~ 0.35 cph), where the forcing is weak and the admittance is estimated from broadband residuals; these regions require 60–120 days to stabilise. This frequency partitioning explains the plateau in Supplementary Figure 4(a): the tidally forced part of the spectrum is learned early and accounts for most of the reducible error, leaving weakly excited components that contribute little to total MAE.

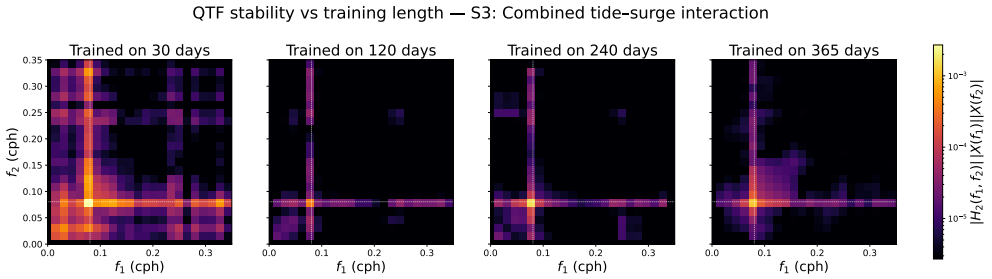

Supplementary Figure 6 Evaluation of weighted quadratic transfer function stability when trained on increasing amounts of data. Applied to synthetic model scenario S3: tide-surge interaction.

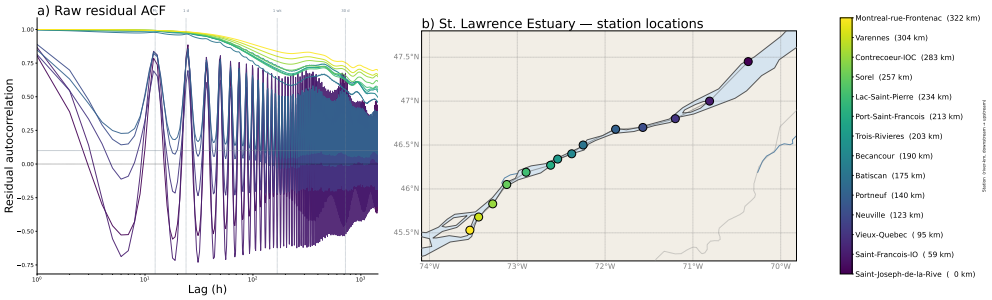

Supplementary Figure 7 Evaluation of residual Autocorrelation Function (ACF) between modeled and observed water levels for Canadian SHOP model.

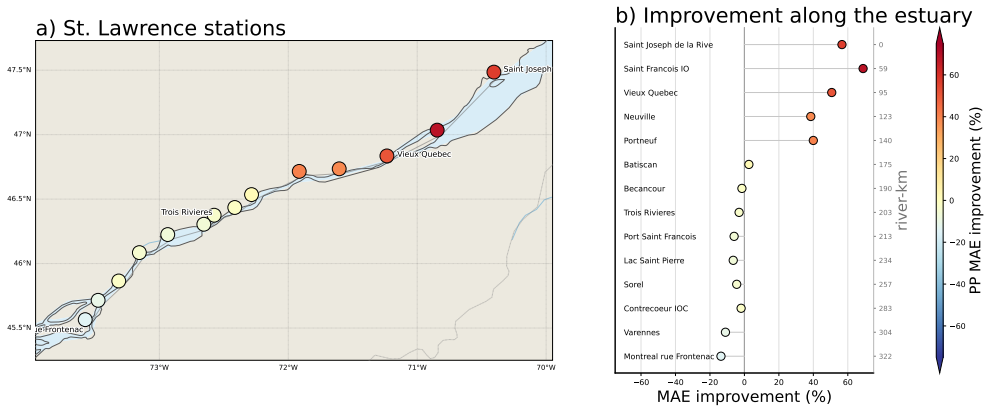

Supplementary Figure 8 Detailed post processed (PP) model results for the Canadian SHOP model in the St. Lawrence Estuary. Panel A shows a map of the sites colored by their relative MAE improvement percentage. The plot on the right visualizes the MAE improvement as a function of distance up river.

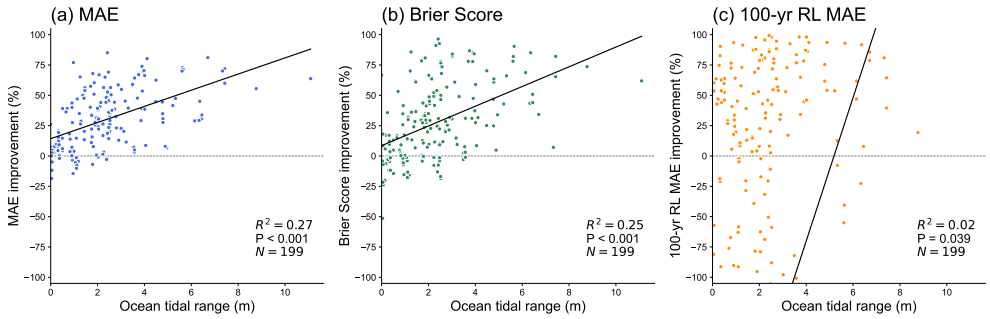

Supplementary Figure 9 Comparisons of postprocessing improvement percentage statistics compared to raw GTSM on GESLA4 gauges. Tidal range estimated using FES2022 tidal atlas [2]. Brier scores correspond to the 99th percentile.

To account for multiple testing across all candidate cuts, p-values were adjusted via a Bonferroni correction ($p_{adj} = \min(1.0, p_{raw} \times m)$). A site was flagged as containing a significant datum shift if the absolute magnitude $|\Delta\mu| \geq 0.05$ m and $p_{adj} \leq 10^{-4}$. If one includes duplicate gauges, this yields 86 gauges; 57 if we account for duplication. As such, we are left with 199 unique sites after removal. All sites removed are listed in Table 1.

4.2 Impact of tides on global postprocessing performance

Observations in the main text suggested that the largest improvements and degradations in postprocessed performance were localized to areas of high and low tides respectively. We consider this relationship by plotting the same statistics as a function of tidal range estimated from the FES2022 tidal atlas [2]. Figure 13 plots this relationship for each metric. The results show that the MAE and 99th percentile Brier score have a statistically significant linear relationship with tidal range (around 25% of the variance is explainable for each). The performance in the 100-yr return period, unsurprisingly, is less correlated to the tidal range. This makes sense as the variability for such return periods is increasingly dominated by meteorological effects. A weak statistically significant relationship $P < 0.05$ is observed which suggests the nonlinear coupling between these processes is important.

4.3 Linear postprocessing of GTSM

We also compare the performance of the linear operator in postprocessing GESLA4 gauges. These results are presented in Figure 10. For the MAE and Brier scores, the two models produce near identical results. The major differences occur in the estimated 100-yr return period. While both models show similar estimated 100-yr return period MAE values (0.20 m vs 0.19 m), the bilinear approach reduces the mean bias from 0.14 m to 0.02 m, while the linear model actually makes it worse increasing to 0.17 m. This suggests that there are nonlinear mechanisms which (i) are not being resolved in GTSM and (ii) are important in driving the total extreme water level.

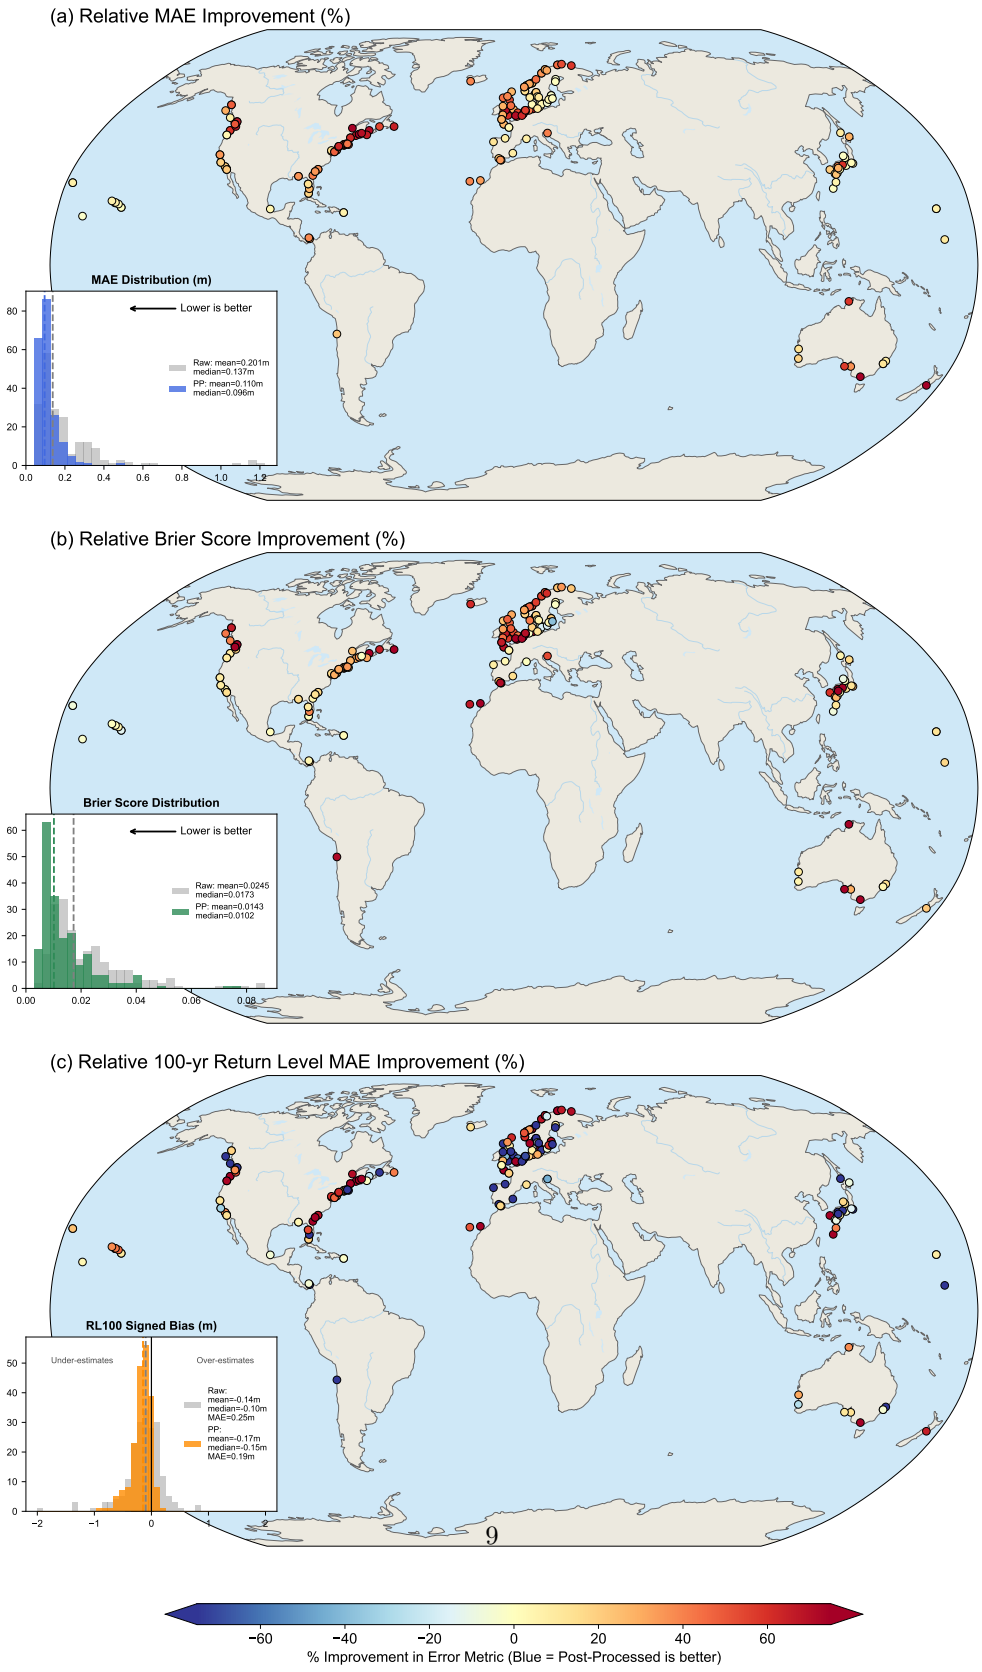

Supplementary Figure 10 Global linear postprocessing of GTSM sea-level reanalysis at GESLA4 tide gauges with 50 years of data (199 sites). Relative improvement in (a) MAE, (b) Brier score, and (c) 100-year return level accuracy under a leave-one-year-out protocol. Insets show the global distributions of each metric before and after postprocessing. Red indicates improvement; Blue indicates degradation. Ocean tidal range is shown in the background and computed from the FES2022 atlas.

5 Comparison of tidal estimators

Lastly, we compare the performance of various estimators that can be used to estimate the coefficients of the linear and bilinear operators. Models are compared across all 43 UK tide gauge network gauges for the same postprocessing task outlined in Section 4 in the main text. We evaluate results when postprocessing AMM15 and AMM75 using three estimators: ordinary least squares (OLS), variational Bayes with ARD priors (VB-ARD — our approach), and a generalised additive model (GAM). We also considered the lag-selection method of Matte et al. [3], but do not report it here as (i) it performed worse than all other estimators in every comparison we ran, and (ii) it was originally designed for connecting adjacent water level stations rather than for the postprocessing task considered here.

Figure 11 shows the results when applying each model across all sites and comparing MAE. For the linear model, the GAM outperforms all other estimators. This is unsurprising: the GAM's smoothness penalty across adjacent lags approximates the physical prior that the impulse response of a linear shallow-water system varies smoothly with τ, and this prior matches the true structure of w_1. However, the GAM fails once bilinear terms are introduced. The GAM framework assumes additive separability of contributions from individual inputs [4], whereas $w_2(\tau_1, \tau_2)$ encodes $(L^{(2)}+1)^2$ pairwise products in which the physically meaningful information resides in specific off-diagonal entries corresponding to cross-frequency couplings. A GAM with bivariate tensor-product smooths would both require pre-specifying which lag pairs interact and, more fundamentally, would impose a smoothness prior on w_2 that is orthogonal to its true structure: sparse, with a small number of concentrated non-zero entries. This is why a prior that directly encodes sparsity is needed once nonlinear interactions enter the model.

Figure 12 compares how adding bilinear terms impacts the postprocessing performance for both the OLS and VB-ARD estimators. Both estimators improve on the linear model on average, but the violin plots show the largest differences between them in the tails of each distribution. The VB-ARD approach has a stunted lower tail compared to OLS, indicating fewer cases in which the addition of bilinear terms degrades performance. It also produces a longer upper tail, with the top end both higher (approximately 96% vs. 77%) and more frequently attained. Both observations are consequences of the same mechanism: the per-coefficient ARD priors shrink coefficients that are not supported by the data while retaining those that are, so redundant or noisy bilinear terms are pruned rather than fit. OLS, by contrast, uses all available interaction terms with equal flexibility, which both allows spurious fits that hurt test performance (degrading the lower tail) and dilutes the contribution of genuinely informative terms (compressing the upper tail).

A further advantage of VB-ARD in this setting is interpretability. The learned sparsity pattern of w_2 is itself physically meaningful — retained non-zero entries correspond to real frequency couplings, and the pattern can be inspected directly rather than recovered via post-hoc feature importance. The posterior precisions α_i also provide uncertainty on each coefficient, enabling well-constrained zero and non-zero entries to be distinguished from ambiguous ones. This makes VB-ARD not only

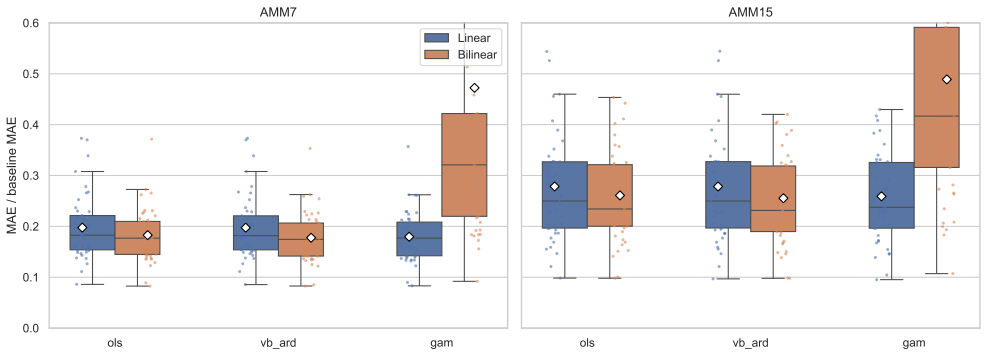

Supplementary Figure 11 Estimator comparisons across UK tide gauge network when applied to both AMM15 and AMM75. Statistics obtained across all UK Tide Gauge Network sites using the same LOYO protocol as in Section 4 of the main text.

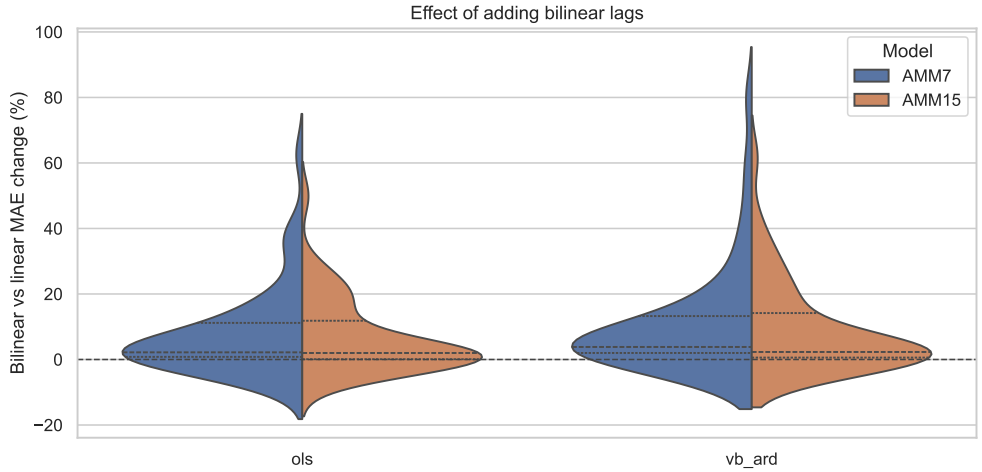

Supplementary Figure 12 Comparison of OLS and VB ARD estimator postprocessing performance improvement when including bilinear terms. Statistics obtained across all UK Tide Gauge Network sites using the same LOYO protocol as in Section 4 of the main text. Positive values corresponding to the bilinear model improving the performance.

the best-performing estimator for the bilinear model, but also the one whose fitted parameters most directly support physical interpretation.

6 Miscellaneous Figures

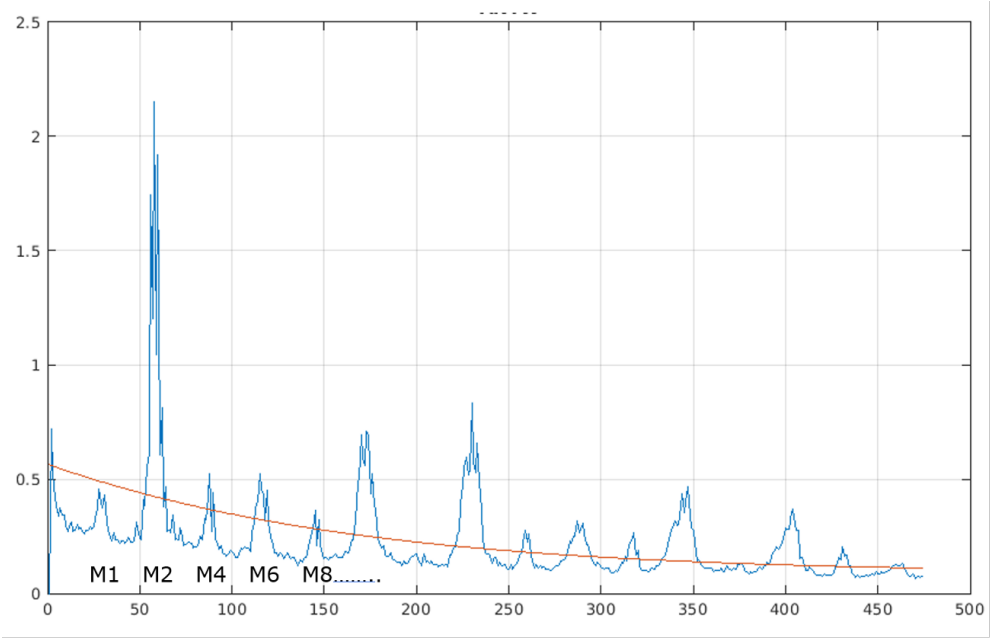

Supplementary Figure 13 Power spectrum of the DCSM7-100 model errors at the Antwerp gauge.

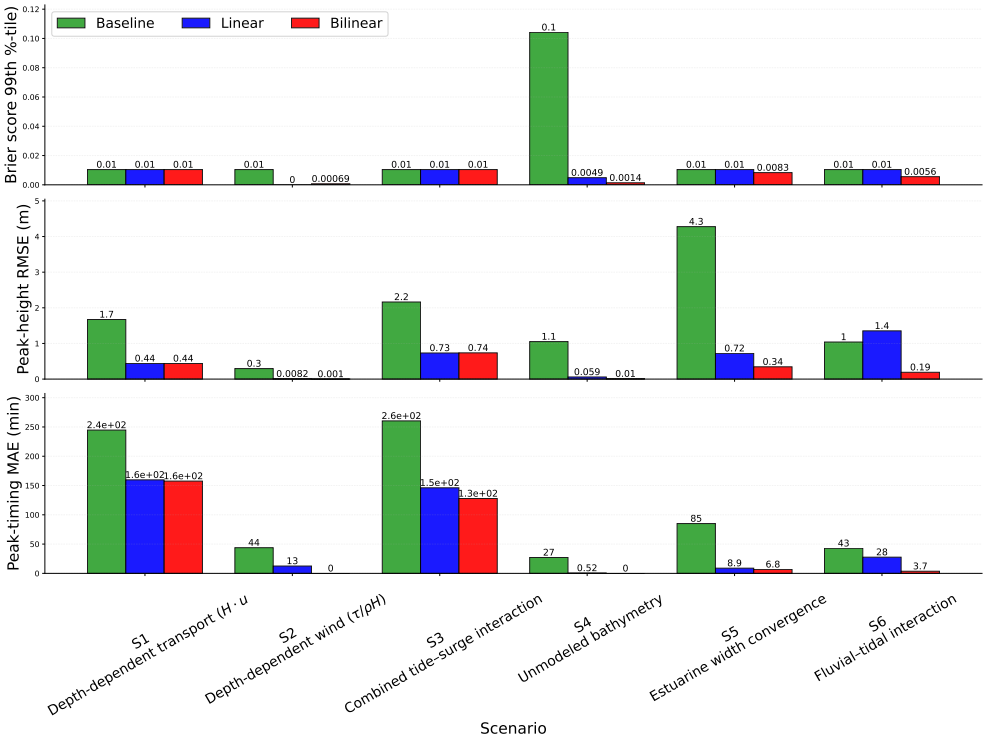

Supplementary Figure 14 Extreme event statistics associated with toy model experiments.

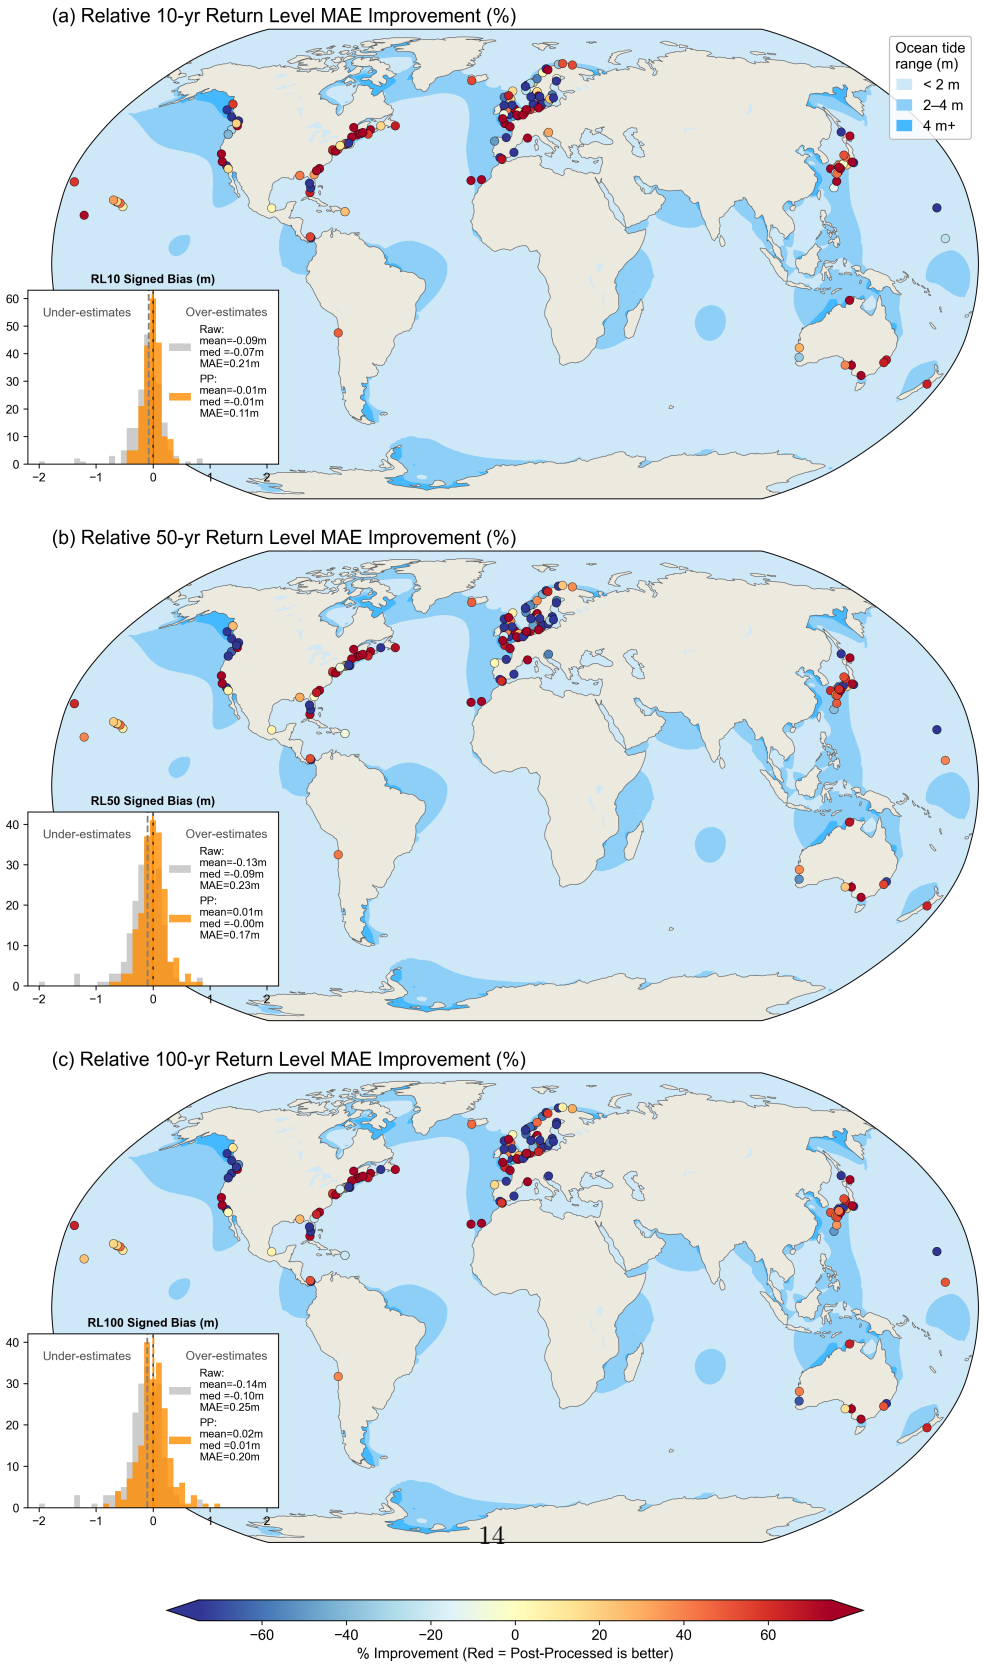

Supplementary Figure 15 Comparison of GTSM and bilinear post-processed return periods.

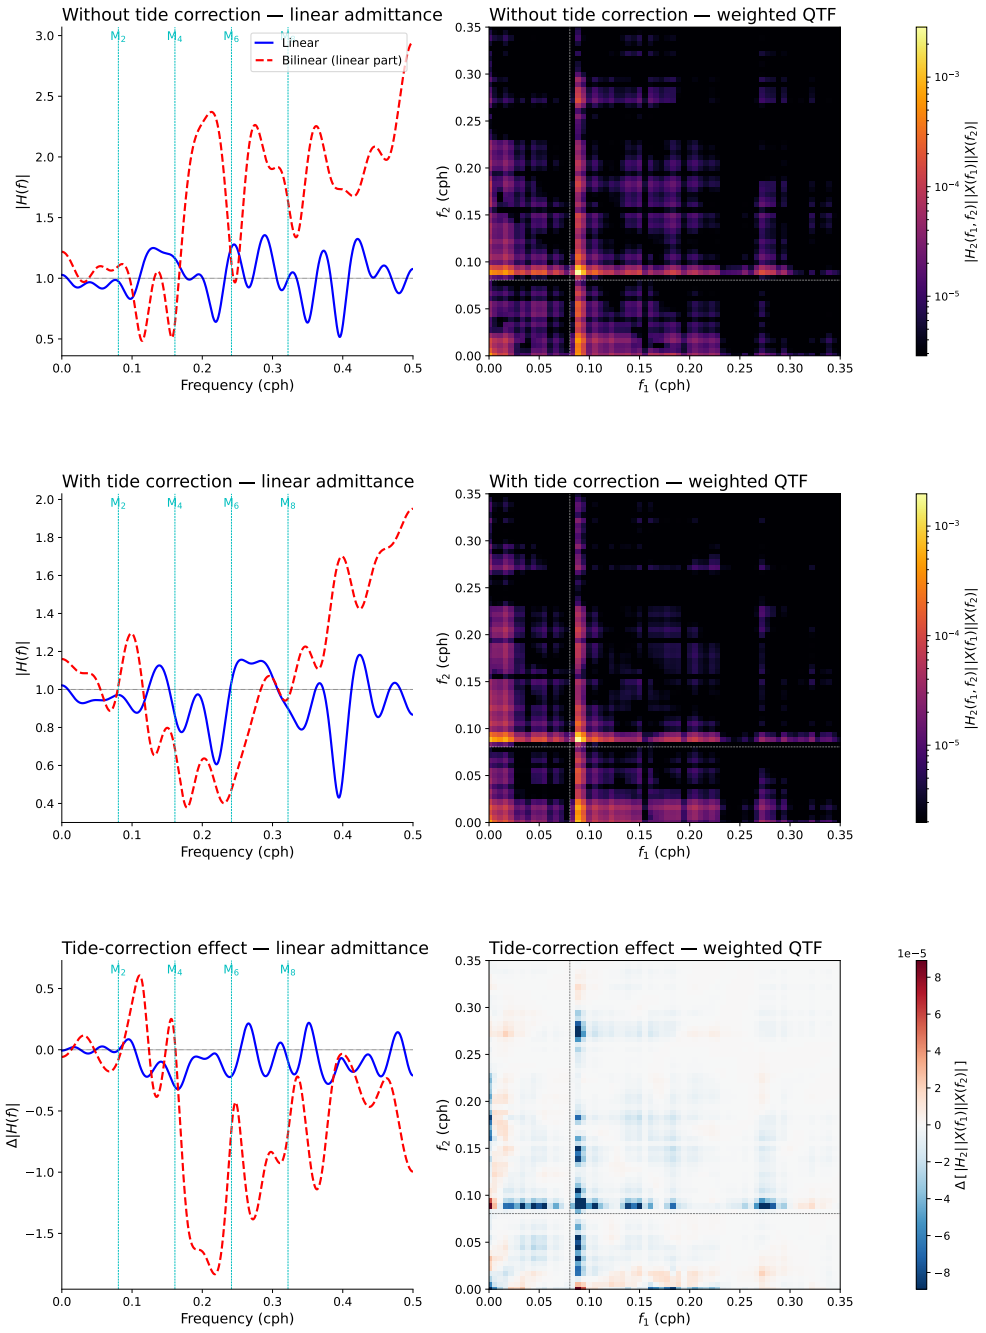

Supplementary Figure 16 Comparison of learned response weights for Steven's Institute operational model with and without the tide correction. Top two rows show the linear admittance and bilinear Weighted Quadratic Transfer Function (WQTF). Bottom row shows the difference between them.

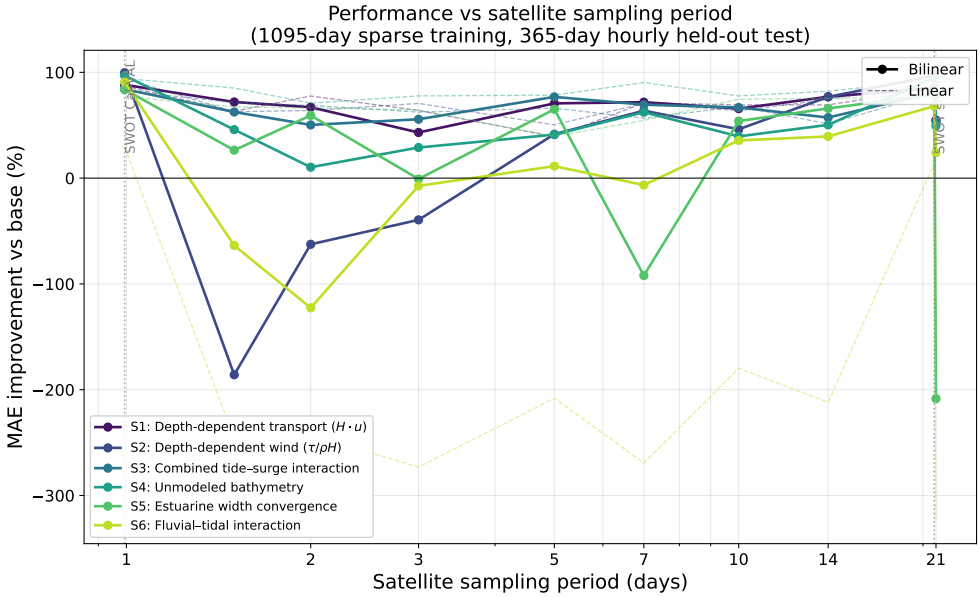

Supplementary Figure 17 Comparison of linear and bilinear postprocessing improvements when applied to increasingly sparse data records for the synthetic case studies. Models are trained on 3-years of data at a given sampling rate ranging from Surface Water Ocean Topography mission's (SWOT) Cal/Val Orbit, to SWOT's maximum science orbit revisit time. Models are then tested out of sample on 60 days of data.

gesla_site	gesla_file	nc_station_id
Galveston_Pier_21	galveston_pier_21-775a-usa-uhslc_rq	38756
Kushimoto	kushimoto-353a-jpn-uhslc_rq	40021
Mera	mera-352a-jpn-uhslc_rq	38882
Tregde	tregde-804a-nor-uhslc_rq	39074
Tofino	tofino-542a-can-uhslc_rq	39068
Neah_Bay_WA	neah_bay_wa-558a-usa-uhslc_rq	38907
Crescent_City_CA	crescent_city_ca-556a-usa-uhslc_rq	38717
Sitka_AK	sitka_ak-559a-usa-uhslc_rq	39038
Kushiro	kushiro-350a-jpn-uhslc_rq	38821
Port_Isabel_TX	port_isabel_tx-772a-usa-uhslc_rq	38949
Apra_Harbor_Guam	apra_harbor_guam-053a-usa-uhslc_rq	39914
Pago_Pago	pago_pago-056a-asm-uhslc_rq	38924
Adak_Alaska	adak_alaska-040a-usa-uhslc_rq	38640
Valparaiso	valparaiso-081a-chl-uhslc_rq	39079
La_Libertad	la_libertad-091a-ecu-uhslc_rq	38826
Hakodate	hakodate-364a-jpn-uhslc_rq	38771
Tumaco	tumaco-303a-col-uhslc_rq	39075
Maisaka	maisaka-356a-jpn-uhslc_rq	38862
Nagasaki	nagasaki-362a-jpn-uhslc_rq	38896
Yakutat_AK	yakutat_ak-570a-usa-uhslc_rq	39100
Ofunato	ofunato-351a-jpn-uhslc_rq	38921
Galveston_Pier_21	galveston_pier_21-775-usa-uhslc_fd	38756
Tregde	tregde-804-nor-uhslc_fd	39074
Tofino	tofino-542-can-uhslc_fd	39068
Neah_Bay_WA	neah_bay_wa-558-usa-uhslc_fd	38907
Crescent_City_CA	crescent_city_ca-556-usa-uhslc_fd	38717
Sitka_AK	sitka_ak-559-usa-uhslc_fd	39038
Kushiro	kushiro-350-jpn-uhslc_fd	38821
Pago_Pago	pago_pago-056-asm-uhslc_fd	38924
Guam	guam-053-usa-uhslc_fd	39914
La_Libertad	la_libertad-091-ecu-uhslc_fd	38826
Valparaiso	valparaiso-081-chl-uhslc_fd	40514
Adak_Alaska	adak_alaska-040-usa-uhslc_fd	38640
Kushimoto	kushimoto-353-jpn-uhslc_fd	40020
Yakutat_AK	yakutat_ak-570-usa-uhslc_fd	39100
Hakodate	hakodate-364-jpn-uhslc_fd	38771
Mera	mera-352-jpn-uhslc_fd	38882
Nishinoomote	nishinoomote-363-jpn-uhslc_fd	40192
Maisaka	maisaka-356-jpn-uhslc_fd	38862
Hachijyojima	hachijyojima-hd11-jpn-jodc_jcg	10210
Miyakejima	miyakejima-hd10-jpn-jodc_jcg	40145
Nishinoomote	nishinoomote-hd21-jpn-jodc_jcg	14811
Omaezaki	omaezaki-ma20-jpn-jodc_jma	40211

Kushimoto	kushimoto-ma26-jpn-jodc_jma	40021
Osaka	osaka-ma30-jpn-jodc_jma	40214
Saigo	saigo-ma54-jpn-jodc_jma	40369
Kushiro	kushiro-ma04-jpn-jodc_jma	38821
Hanasaki	hanasaki-ma03-jpn-jodc_jma	39923
Hakodate	hakodate-ma06-jpn-jodc_jma	38771
Ofunato	ofunato-ma10-jpn-jodc_jma	38921
Onahama	onahama-ma12-jpn-jodc_jma	40212
Shimizuminato	shimizuminato-ma19-jpn-jodc_jma	40420
Maisaka	maisaka-ma21-jpn-jodc_jma	38862
Uno	uno-ma33-jpn-jodc_jma	40505
Mera	mera-ma14-jpn-jodc_jma	38882
Kobe	kobe-ma31-jpn-jodc_jma	40009
Nagasaki	nagasaki-ma49-jpn-jodc_jma	38896
Hosojima	hosojima-gs02-jpn-jodc_giaj	39951
Aburatsubo	aburatsubo-gs01-jpn-jodc_giaj	14884
Nezugaseki	nezugaseki-gs07-jpn-jodc_giaj	10270
Tajiri	tajiri-gs10-jpn-jodc_giaj	15016
Ilfracombe	ilfracombe-ilf-gbr-bodc	39957
Cromer	cromer-cro-gbr-bodc	39412
Galveston_Pier_21	galveston_pier_21-8771450-usa-noaa	38756
Sewells_Point	sewells_point-8638610-usa-noaa	40416
Neah_Bay	neah_bay-9443090-usa-noaa	38907
Crescent_City	crescent_city-9419750-usa-noaa	38717
Sitka	sitka-9451600-usa-noaa	39038
Port_Isabel	port_isabel-8779770-usa-noaa	38949
Adak_Island	adak_island-9461380-usa-noaa	38640
Yakutat	yakutat-9453220-usa-noaa	39100
Tofino_BC	tofino_bc-8615-can-meds	39068
Churchill	churchill-5010-can-meds	39780
Catania	catania-cat-ita-ispra	39753
Livorno	livorno-liv-ita-ispra	40064
Civitavecchia	civitavecchia-civ-ita-ispra	39785
Messina	messina-mes-ita-ispra	40133
Ravenna	ravenna-rav-ita-ispra	40343
Porto_Empedocle	porto_empedocle-por-ita-ispra	40297
Stavanger	stavanger-svg-nor-nhs	39362
Tregde	tregde-trg-nor-nhs	40495
StavangerTG	stavangertg-sta-nor-cmems	39362
NarvikTG	narviktg-nar-nor-cmems	39607

Supplementary Table 1: GESLA4 gauges [5] identified with $\geq$ 0.1m datum shifts over 50-year analysis.

References

[1] Higgins, J. R. *Sampling theory in Fourier and signal analysis: foundations* (OUP Oxford, 1996).

[2] CNES. Fes2022 (finite element solution) tidal model (version 2024). Data set (2024). URL https://doi.org/10.24400/527896/A01-2024.004. The FES2022 Tide product was funded by CNES, produced by LEGOS, NOVELTIS and CLS and made freely available by AVISO.

[3] Matte, P. & Innocenti, S. Filling gaps and extending inland tide gauge records to support accurate long-term and extreme water level analyses. *Continental Shelf Research* 105668 (2026).

[4] Hastie, T. J. Generalized additive models. *Statistical models in S* 249–307 (2017).

[5] Haigh, I. D. *et al.* Gesla version 4: Further enhancements to the global high-frequency sea-level dataset (2026). In preparation.

[45] Morin, J., Leclerc, M., Secretan, Y. & Boudreau, P. Integrated two-dimensional macrophytes-hydrodynamic modeling. *Journal of Hydraulic Research* **38**, 163–172 (2000).

[46] Matte, P., Secretan, Y. & Morin, J. Quantifying lateral and intratidal variability in water level and velocity in a tide-dominated river using combined RTK GPS and ADCP measurements. *Limnology and Oceanography: Methods* **12**, 281–302 (2014).

[47] Matte, P., Secretan, Y. & Morin, J. Hydrodynamic Modeling of the St. Lawrence Fluvial Estuary. I: Model Setup, Calibration, and Validation. *Journal of Waterway, Port, Coastal, and Ocean Engineering* **143**, 4017010 (2017).

[48] Matte, P., Secretan, Y. & Morin, J. Hydrodynamic Modeling of the St. Lawrence Fluvial Estuary. II: Reproduction of Spatial and Temporal Patterns. *Journal of Waterway, Port, Coastal, and Ocean Engineering* **143**, 4017011 (2017).

[49] Matte, P., Secretan, Y. & Morin, J. Erratum for "Hydrodynamic Modeling of the St. Lawrence Fluvial Estuary. I: Model Setup, Calibration, and Validation" by Pascal Matte, Yves Secretan, and Jean Morin. *Journal of Waterway, Port, Coastal, and Ocean Engineering* **144**, 8218001 (2018).

[50] Warren, I. & Bach, H. Mike 21: a modelling system for estuaries, coastal waters and seas. *Environmental software* **7**, 229–240 (1992).

[51] Kernkamp, H. W., Van Dam, A., Stelling, G. S. & de Goede, E. D. Efficient scheme for the shallow water equations on unstructured grids with application to the continental shelf. *Ocean Dynamics* **61**, 1175–1188 (2011).

[52] Virtanen, P. *et al.* Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods* **17**, 261–272 (2020).

[53] Muis, S., Aleksandrova, N., Veenstra, J. & Gwee, R. Gtsm-era5-e dataset - data underlying the paper "global dataset of storm surges and extreme sea levels for 1950-2024 based on the era5 climate reanalysis" (2025). URL https://doi.org/10.5281/zenodo.14671593.

[54] Muis, S. *et al.* Global sea level change time series from 1950 to 2050 derived from reanalysis and high resolution CMIP6 climate projections (2022). Accessed on [April 2026].