JGR | Machine Learning and Computation

RESEARCH ARTICLE
10.1029/2024JH000523

Computationally Efficient Hybrid Downscaling of Surf Zone Hydrodynamics: Methodology and Evaluation

E. R. Echevarria[1], S. Contardo[2], B. Pérez-Díaz[3], R. K. Hoeke[4], B. Leighton[5], C. Trenham[6], L. Cagigal[3], and F. J. Méndez[3]

[1]Environment, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Hobart, TAS, Australia, [2]Environment, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Crawley, WA, Australia, [3]Departamento de Ciencias y Técnicas del Agua y del Medio Ambiente, Geomatics and Ocean Engineering Group, E.T.S.I. C.C.P, Universidad de Cantabria, Santander, Spain, [4]Environment, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Aspendale, VIC, Australia, [5]Environment, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Clayton, VIC, Australia, [6]Environment, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Black Mountain, ACT, Australia

Key Points:

- A hybrid XBeach model integrates numerical simulations with machine learning for efficient coastal predictions
- The surrogate model efficiently and accurately predicts nearshore wave and hydrodynamic parameters
- Data reduction via Empirical Orthogonal Functions provides enhanced efficiency with minimal accuracy loss

Supporting Information:

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

Correspondence to:

E. R. Echevarria,
emilio.echevarria@csiro.au

Citation:

Echevarria, E. R., Contardo, S., Pérez-Díaz, B., Hoeke, R. K., Leighton, B., Trenham, C., et al. (2025). Computationally efficient hybrid downscaling of surf zone hydrodynamics: Methodology and evaluation. *Journal of Geophysical Research: Machine Learning and Computation*, *2*, e2024JH000523. https://doi.org/10.1029/2024JH000523

Received 14 NOV 2024
Accepted 5 MAY 2025

Author Contributions:

Conceptualization: E. R. Echevarria, S. Contardo, B. Pérez-Díaz, R. K. Hoeke, B. Leighton, C. Trenham, L. Cagigal, F. J. Méndez
Investigation: E. R. Echevarria, S. Contardo, B. Pérez-Díaz, R. K. Hoeke, B. Leighton, C. Trenham, L. Cagigal, F. J. Méndez
Methodology: E. R. Echevarria, S. Contardo, B. Pérez-Díaz, R. K. Hoeke, B. Leighton, C. Trenham, L. Cagigal, F. J. Méndez

Abstract We present a hybrid surf-zone model that combines numerical simulations and statistical/machine learning techniques, enabling accurate calculations of nearshore wave and hydrodynamic parameters with high computational efficiency. The approach involves defining representative forcing conditions, carrying out numerical model (XBeach) simulations for these cases, and training machine learning models capable of predicting selected model output variables. Data decomposition via Empirical Orthogonal Function analysis further simplifies the process, reducing the output data dimensionality, with minimal accuracy loss (with exception of certain wetting-drying processes). Three machine learning approaches of increasing complexity are compared: a multi-variate linear regression (LR), a Radial Basis Functions (RBF) interpolator and a Deep Neural Network (DNN). The LR model fails to account for the complex non-linearities in coastal wave dynamics, which warrants the use of more complex machine learning techniques. Both the RBF interpolator and the DNN models demonstrate high levels of accuracy in the prediction of short wave heights, mean wavelength, and depth-averaged currents, with slightly lower accuracy for long (infragravity) wave heights and fraction of breaking waves. The proposed surrogate model thus offers an efficient alternative to computationally expensive numerical model simulations, enabling rapid and reliable long-period deterministic simulations (multi-decadal hindcasts) and/or multi-ensemble probabilistic scenario simulations of nearshore hydrodynamic conditions. We provide a comprehensive description of the implementation details and assess the surrogate model's performance in representing various wave and hydrodynamic parameters. We discuss potential use cases and limitations, noting that this hybrid modeling technique can be adapted for use with other numerical models in various settings.

Plain Language Summary This study presents an efficient and accurate approach to modeling nearshore wave conditions, water levels and currents, by combining numerical model simulations with machine learning techniques. Traditional numerical models are computationally intensive, while the hybrid approach presented here is a much faster alternative. The study provides a comprehensive, step-by-step guide for building this surrogate model, which involves defining a set of representative conditions and training a machine learning model to capture the relationship between input forcings and output variables (gridded fields of selected parameters). To improve efficiency, an Empirical Orthogonal Function analysis is applied as a data reduction step, simplifying the process without substantially compromising accuracy. This surrogate model demonstrates high precision, and can be used to calculate long-term climatologies of nearshore wave and hydrodynamic conditions (something that would be prohibitively expensive with the traditional numerical model), or the almost instantaneous production of nearshore forecasts.

1. Introduction

Coastal and nearshore areas are complex, highly energetic and dynamic environments. These areas are home to an array of ecosystems, human settlements and infrastructure, and they are exposed to numerous threats, including erosion, storm surges, inundation and extreme wave conditions (McGranahan et al., 2007). Additionally, these areas are critical for ocean safety, operational activities, and economic pursuits including fishing and tourism, coastal defense, recreational use, and marine conservation efforts. Climate change exacerbates the vulnerability of

Software: E. R. Echevarria, S. Contardo, B. Pérez-Díaz, R. K. Hoeke, B. Leighton, C. Trenham, L. Cagigal, F. J. Méndez
Validation: E. R. Echevarria, S. Contardo, R. K. Hoeke, B. Leighton, C. Trenham
Writing – original draft: E. R. Echevarria, S. Contardo, R. K. Hoeke, C. Trenham
Writing – review & editing: E. R. Echevarria, S. Contardo, B. Pérez-Díaz, R. K. Hoeke, C. Trenham, L. Cagigal, F. J. Méndez

coastal areas, with an increase in the sea level and the frequency and severity of extreme weather events. Therefore, understanding coastal and nearshore processes is fundamental to addressing these challenges effectively, enhancing resilience and ensuring the sustainable development of coastal regions.

In this context, numerical models are an invaluable tool for better understanding the dynamics of coastal systems and for quantifying and predicting oceanographic conditions at the nearshore. XBeach (Roelvink et al., 2009) is an open-source, process-based morphodynamic numerical model widely used for analyzing nearshore wave dynamics, currents, sediment transport and sea-level conditions, and their impacts on the coast. For incident waves, XBeach surf-beat mode solves the short-wave action balance on the scale of wave groups, resolving the transformation processes waves undergo as they approach shallow waters (refraction, shoaling, breaking). In addition, long (infragravity) wave motions and mean flow can be fully resolved using the nonlinear shallow water equations (Phillips, 1977). Thus, wave-driven currents, wave setup and swash processes are also included. Importantly, XBeach is capable of modeling sediment transport processes and morphological changes of the coast and nearshore areas. XBeach has been successfully implemented to quantify the impacts of tropical cyclones on inundation levels (e.g., Tu'uholoaki et al., 2023), to better understand the dynamics of infragravity wave heights (e.g., Pomeroy et al., 2012), and to study the morphological response of coastal settings (e.g., Smallegan et al., 2016), among many other applications. However, increased accuracy and complexity in numerical models entails elevated computational costs. Multi-decadal and/or multi-ensemble simulations of wave processes at the nearshore, required for probabilistic climate and scenario prediction, are prohibitively computationally expensive with XBeach. Although GPU-based alternatives are being developed (Rautenbach et al., 2022) which reduce these costs, they remain large and the alternatives to date lack the full functionality of XBeach. Coastal observing networks, while providing highly accurate data, often lack the spatial resolution necessary to capture the dynamics of nearshore processes. Similarly, weather forecasts and climate models typically do not adequately resolve the complexity of these coastal phenomena. Therefore, statistical methods, such as Extreme Value Analysis (EVA), that rely on these data sets often carry substantial uncertainties and have limited applicability at local scales (Hinkel et al., 2021). These limitations underscore the need for more sophisticated modeling approaches that can better account for the complexities of nearshore environments.

In this paper we investigate a hybrid approach, combining XBeach numerical simulations and machine learning (ML) techniques, aimed at significantly reducing the high computational costs of the numerical model for such applications whilst retaining accuracy. Importantly, this methodology predicts the spatial (gridded) output fields of XBeach, rather than predefined coastal point locations, which is crucial for comprehensive coastal management and planning. Spatial wave, current and water level predictions are particularly important for applications within the surf-zone. Comprehensive predictions enable better planning and preparedness for coastal defense, infrastructure development, and environmental conservation, providing a holistic view of potential impacts across varied scenarios.

Hybrid methodologies have emerged to bridge the gap between computational efficiency and physical accuracy, combining the strengths of dynamical (numerical) models and statistical tools (e.g., Nieves et al., 2021; Ricondo et al., 2023; Serafin et al., 2019; Tausía et al., 2023; Wang et al., 2024; Wang et al., 2025). Hybrid models leverage the predictive power of ML models trained on outputs from process-based models, enabling rapid, data-driven predictions without sacrificing accuracy. Unlike purely statistical models, hybrid approaches better retain both the complexity and the geophysical constraints of dynamic model simulations, whilst significantly reducing computational demands of dynamic model-only solutions. The goal of a supervised ML algorithm is to relate inputs and corresponding outputs by emulating or approximating the sometimes intricate and complex non-linear relationships between them. This is particularly true in the context of coastal wave dynamics, where non-linear effects such as bottom friction, wave breaking, and other wave transformations play a significant role. There is an immense variety of approaches, techniques and methodologies aimed at improving the understanding of coastal and nearshore oceanographic conditions, as well as predicting them. The reader is referred to review papers on this topic (e.g., Goldstein et al., 2019; Kim & Lee, 2022). In this paper we closely follow the methods developed by Zornoza-Aguado et al. (2024) and expand them to include evaluation of both simpler (multivariate linear regression, LR) and more complex (deep neural network, DNN) approaches in addition to the radial basis function (RBF) algorithm. The incorporation of DNNs offers greater flexibility to capture complex, non-linear relationships in the data by allowing adjustments to the network's architecture and components to optimize performance, potentially improving predictive accuracy. This contrasts with LR, which only captures the simpler correlations of dependent variables. We also consider potential limitations of the empirical orthogonal function (EOF) data

10.1029/2024JH000523

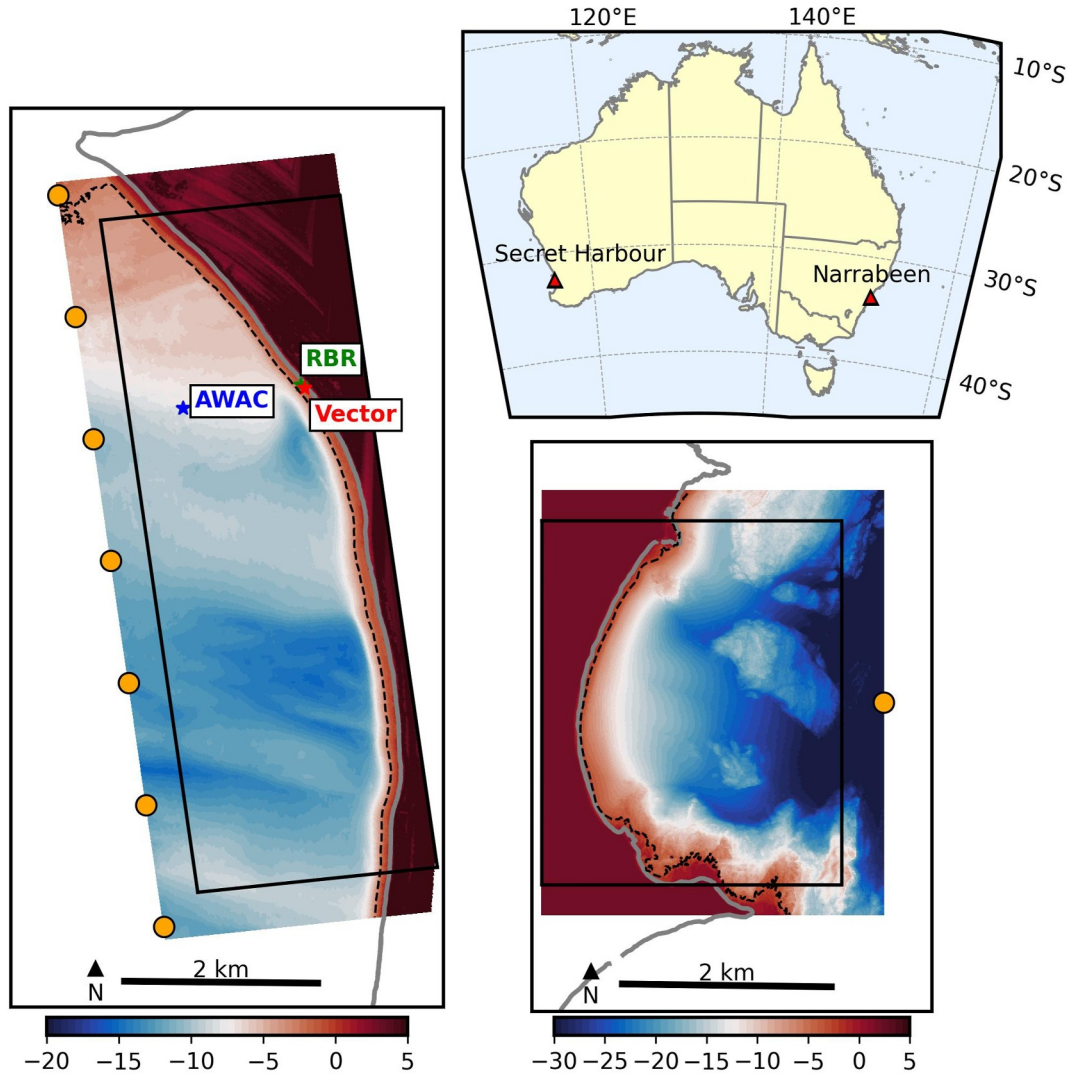

Figure 1. XBeach-SB model domains and location maps. Upper right panel: location map of the two model domains in Australia. Left panel: the Secret Harbour, Western Australia domain; location of in situ observations discussed in the text are indicated. Lower right panel: the Narrabeen, New South Wales domain. In both the (left and lower right) model domain plots, the colored bathymetry and topography indicate the XBeach-SB total domain extent, solid black lines indicate the areas used for surrogate model training and numerical/surrogate model comparisons and the location of offshore open boundary condition points are indicated with orange circles; for reference, the dotted black line indicates the −3 m contour and the light gray line indicates the shoreline, both at mean sea level.

reduction step, compare both fully dynamic and hybrid results to observations and consider a wide range of hydrodynamic output variables within the final hybrid model.

We apply and evaluate these hybrid methods at two different domains in Australia which experience distinct wave climates and tidal/sea level variability regimes: one in Secret Harbour, Western Australia and the other in Narrabeen, New South Wales (Figure 1). These locations and respective model configurations are described in Section 2. In Section 3 we present the methods used to implement ML for prediction of wave parameters. Results of these approaches are given in Section 4. We also provide an Appendix describing the deep neural network design used in this study.

2. Forcing Data and XBeach Domain Configurations

At both study sites, XBeach is run in surf-beat mode (XBeach-SB hereafter), therefore resolving the infragravity wave motions and transformations. Water level (*wl*) boundary conditions for both location domains

(Figure 1) are provided by the linear addition of TPXO9.2 (Egbert & Erofeeva, 2002) predicted tides and monthly sea surface height from the ECMWF-ORAS5 global ocean reanalysis (Copernicus Climate Change Service, https://doi.org/10.24381/cds.67e8eeb7), providing ∼40-year (1980–2024) hourly time series. Due to the complexity of offshore reefs and bathymetry in the vicinity of Secret Harbour, wave boundary conditions for this domain are drawn from a regional 500 m spatial resolution SWAN wave model downscale of the CAWCR Global Wave Hindcast's regional Australia 4-arcminute domain (Trenham et al., 2024); whereas at the Narrabeen location, which has relatively simpler and deeper offshore bathymetry, wave boundary conditions are drawn directly from the CAWCR Wave Hindcast Australia 4-arcminute domain output (Smith et al., 2021). At both locations, XBeach-SB sediment transport processes are turned off, as the focus of this study is on nearshore hydrodynamics. At each of the boundary points, integrated wave parameters (*Hs*: significant wave height; *Tp*: peak wave period; *Dp*: peak wave direction; *spr*: directional wave spreading) are used to generate JONSWAP spectra that provide the wave forcing conditions to the model.

Bathymetry and nearshore topography for both the Secret Harbour and Narrabeen model domains are based on 10 m gridded digital elevation models (DEMs) derived from Fugro LADS Corporation airborne LiDAR surveys performed in 2009 and 2018, respectively. Secret Harbour region's DEM was provided by the Western Australian government's Department of Transport (https://www.transport.wa.gov.au/imarine/marine-geographic-data.asp); the Narrabeen region's DEM was provided by (State Government of NSW and NSW Department of Climate Change et al., 2023). The next two sub-sections describe these two study sites/model domains in greater detail.

2.1. Secret Harbour Beach, Western Australia

Secret Harbour is located in south-western Western Australia. It experiences a diurnal microtidal regime (mean tidal range <0.5 m) and a wave climate consisting of swell waves arriving remotely from the Southern Ocean (predominantly during winter) and locally generated wind waves due to strong sea breezes predominantly during summer; both wave sources approach the coastline from the south to south-west and it is further partly sheltered from open ocean swell by a mostly continuous offshore reef located ∼6 km from the shore with a minimum depth of ∼5 m (Contardo et al., 2019; Segura et al., 2018). The model domain is defined with a rotated regular grid spanning approximately 9 km in the alongshore direction with a resolution of 15 m, and approximately 3 km with a resolution of 10 m in the cross-shore direction.

At this study site, in situ observations from previous studies are available (e.g., Contardo & Symonds, 2013). Instruments used in this study (located as shown in the left panel of Figure 1) include:

- A Nortek acoustic wave and current (AWAC) profiler (at ∼7.16 m depth); the relevant observations used here are hourly measurements of integrated wave parameters (especially *Hs*), derived from the instruments acoustic surface tracking.
- An RBR 1 Hz pressure sensor (at ∼0.9 m depth); observations were averaged at 3,600 s intervals to provide hourly *wl* (including local wave setup) observations.
- A Nortek vector velocimeter with a 1 Hz pressure sensor (at ∼1.6 m depth) was used to calculate the infra-gravity wave height (H_{ig}), derived from the instantaneous water level elevation measurements.

These observations are used to assess the accuracy of the XBeach-SB configuration (and later used to compare to that of the surrogate model).

2.2. Narrabeen Beach, New South Wales

Approximately 3,300 km away, Narrabeen Beach in central New South Wales is located on the opposite side of Australia from Secret Harbour (Figure 1); it experiences semi-diurnal tides, and while it is still classified as micro-tidal (mean spring range of 1.3 m, e.g. see Turner et al., 2016), the range is roughly double that of Secret Harbour. Local wave climate is dominated by a mix of sources: trade-wind swell associated with the subtropical high, Tasman Sea mid-latitude cyclones and Australian East Coast Lows (Mortlock & Goodwin, 2015); wave incidence tends to be much less oblique than at Secret Harbour, but is also considerably more energetic and variable. The model domain consists of a regular grid which extends approximately 4.9 km in the alongshore direction and approximately 3.3 km in the cross-shore; the lateral boundaries coincide with rocky headlands and adjacent reefs that delineate Narrabeen Beach. A few rocky reefs exist near the center of the domain, but few reefs with significant impacts on wave propagation occur directly offshore from the domain. Spatial resolution in both the

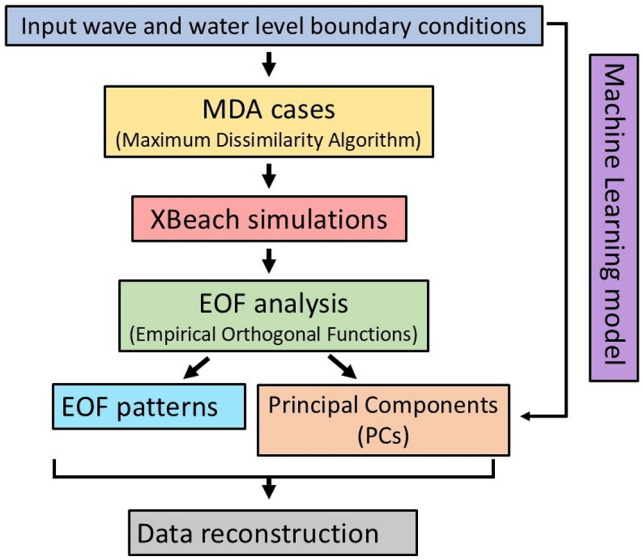

Figure 2. Schematic representation of the hybrid numerical model/ML methodology proposed in this study.

alongshore and cross-shore directions is 10 m. Wave boundary conditions are drawn from the CAWCR Wave Hindcast at a point in the middle of the offshore boundary. Integrated wave parameters are used to generate JONSWAP spectra to provide wave boundary conditions to the model.

3. Hybrid Methodology

A schematic of the methodology used in this study is presented in Figure 2. The main objective is to predict or "emulate" the output of XBeach-SB simulations from the input parameters used to run those simulations leveraging ML techniques. The initial step is related to the model training/learning. For the ML model to accurately represent a range of potential scenarios, it needs to be trained (and tested) using a comprehensive data set encompassing a wide range of possible wave and sea level conditions. An effective way of achieving a set of representative conditions of a given N-dimensional data set is through use of the Maximum Dissimilarity Algorithm (MDA, see Camus et al. (2011) for a detailed description of this methodology). MDA extracts a set of the most dissimilar cases in a certain data set, defined as those that are the farthest apart based on the Euclidian distance between the cases. In our case, the data set is comprised of the parameters that define the wave boundary conditions (*Hs*, *Tp*, *Dp*, and *spr*, described in Sections 2.1 and 2.2 for each respective study site) and the water level elevations (*wl*, accounting for tides and background sea level variability). The MDA is performed on this 5-dimensional data set (*Hs*, *Tp*, *Dp*, *spr*, *wl*) to extract the 1000 most dissimilar cases, thus ensuring that the resulting XBeach-SB simulations will encompass a wide range of possible scenarios covering this parameter space. The training data set is generated using data from 1980 to 2015, while the testing data is generated running the MDA analysis on the last 5 years of hindcast data (2015–2019). For the Secret Harbour case, the MDA is applied to data from the central offshore boundary point among seven total boundary points. This yields 1,000 representative training cases based on conditions at this central point. To simulate the full boundary conditions, we record the times in the hindcast corresponding to these representative cases and then extract data from all seven boundary points at those specific times.

The results of the application of the MDA are shown in Figure 3, where the gray dots show the entire hindcast data set, the blue dots denote the 1,000 MDA cases used for training, and the red dots show an independent data set of 1,000 MDA cases that will be used for testing the performance of the surrogate model in representing unseen data and prevent over-fitting. The MDA is a particularly useful technique for this hybrid methodology, as it covers the diversity in the input parameter space with a reduced number of representative cases. However, the MDA approach is inherently limited to the input data set on which it was performed, meaning that it may not account for potential future scenarios, such as those induced by climate change or tropical cyclones, which could introduce new, unseen conditions. To account for the projected impacts of climate change, Anderson et al. (2021) devised an approach in which the MDA is performed on data generated by a stochastic climate model emulator, thus producing plausible synthetic time series of future scenarios. This approach enables surrogate models to capture potential long-term variations in wave and water level conditions, improving their applicability for future climate assessments. Additionally, in cases where detailed historical data are not available, alternative approaches for generating training/testing data sets, such as the Latin Hypercube Sampling approach (McKay et al., 2000) can be implemented instead (see, for example, Zornoza-Aguado et al. (2024) and Ricondo et al. (2024)).

Next, the numerical model is run for these MDA scenarios (in this case, for example, the XBeach-SB model for the Secret Harbour domain). The model is run for each of the 2,000 MDA cases (encompassing both training and testing data sets). The simulations are run separately, for a period of 3 hr, with the initial 2 hr considered as spin-up and thus excluded from the analysis. The simulations utilize the previously described hourly wave boundary conditions and water levels as forcing. The timestep used for defining the wave energy and long wave flux at the offshore boundary (keyword *dtbc* in XBeach) was set to 2 s. The CFL time step was set to 0.7.

The following step in this methodology involves a dimensionality reduction of the XBeach-SB output fields by means of a decomposition with Empirical Orthogonal Functions (EOF, e.g. see Björnsson & Venegas, 1997).

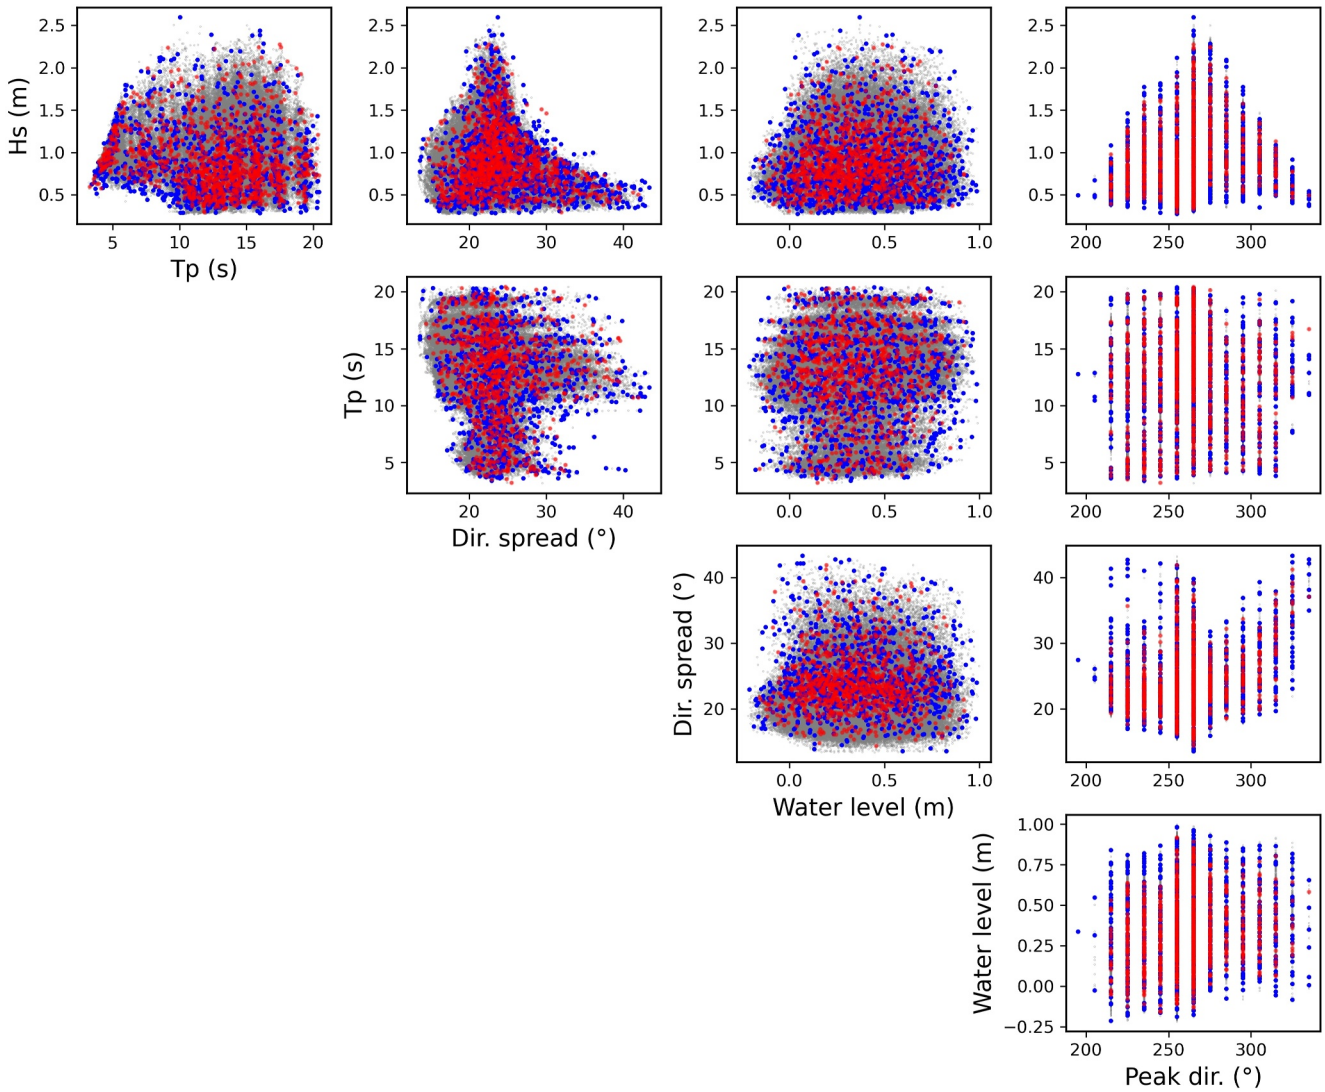

Figure 3. Scatter plots of XBeach-SB forcing parameters (significant wave height, peak period, directional spread, water level and peak direction) for the Secret Harbour domain. The gray points show the entire hindcast data set (from 1980 till 2019). The blue points show the 1,000 MDA cases selected for training the ML model (MDA applied to hindcast data from 1980 to 2015), and the red dots are an independent data set of 1,000 MDA cases used for testing (MDA using data from 2015 to 2019).

Given a gridded (longitude/latitude) time-varying field of a scalar quantity (in this case XBeach-SB outputs, such as the wave fields), the EOF analysis identifies and isolates the dominant patterns or modes of variability, decomposing the data in what will be termed here EOF patterns (these are the eigenvectors of the data covariance matrix, and represent the main spatial patterns of variability) and Principal Components (PCs), which are projections of the data on the EOF patterns, and represent the temporal variability of these. The original data can then be reconstructed using the following equation:

$$\text{reconstructed data} = \Sigma(\text{EOF}_i(x,y) \times \text{PC}_i(t)) \qquad i = 1, \ldots, N \tag{1}$$

Here, $\text{EOF}_i(x,y)$ represents the EOF patterns and $\text{PC}_i(t)$ the PCs. It should be noted that there is no explicit time dimension in the MDA data sets, therefore in this context t represents each independent MDA case (1, 2, …, 1,000). The accuracy of the reconstruction is determined by the number of EOFs retained (N). The methodology outlined in this paper consists of predicting the PCs that result from this decomposition, rather than the raw, temporally varying gridded XBeach-SB outputs. First, the training data set is decomposed into EOF patterns and PCs. For the testing data set, the same EOF patterns are used, and only the PCs are calculated by projecting the

10.1029/2024JH000523

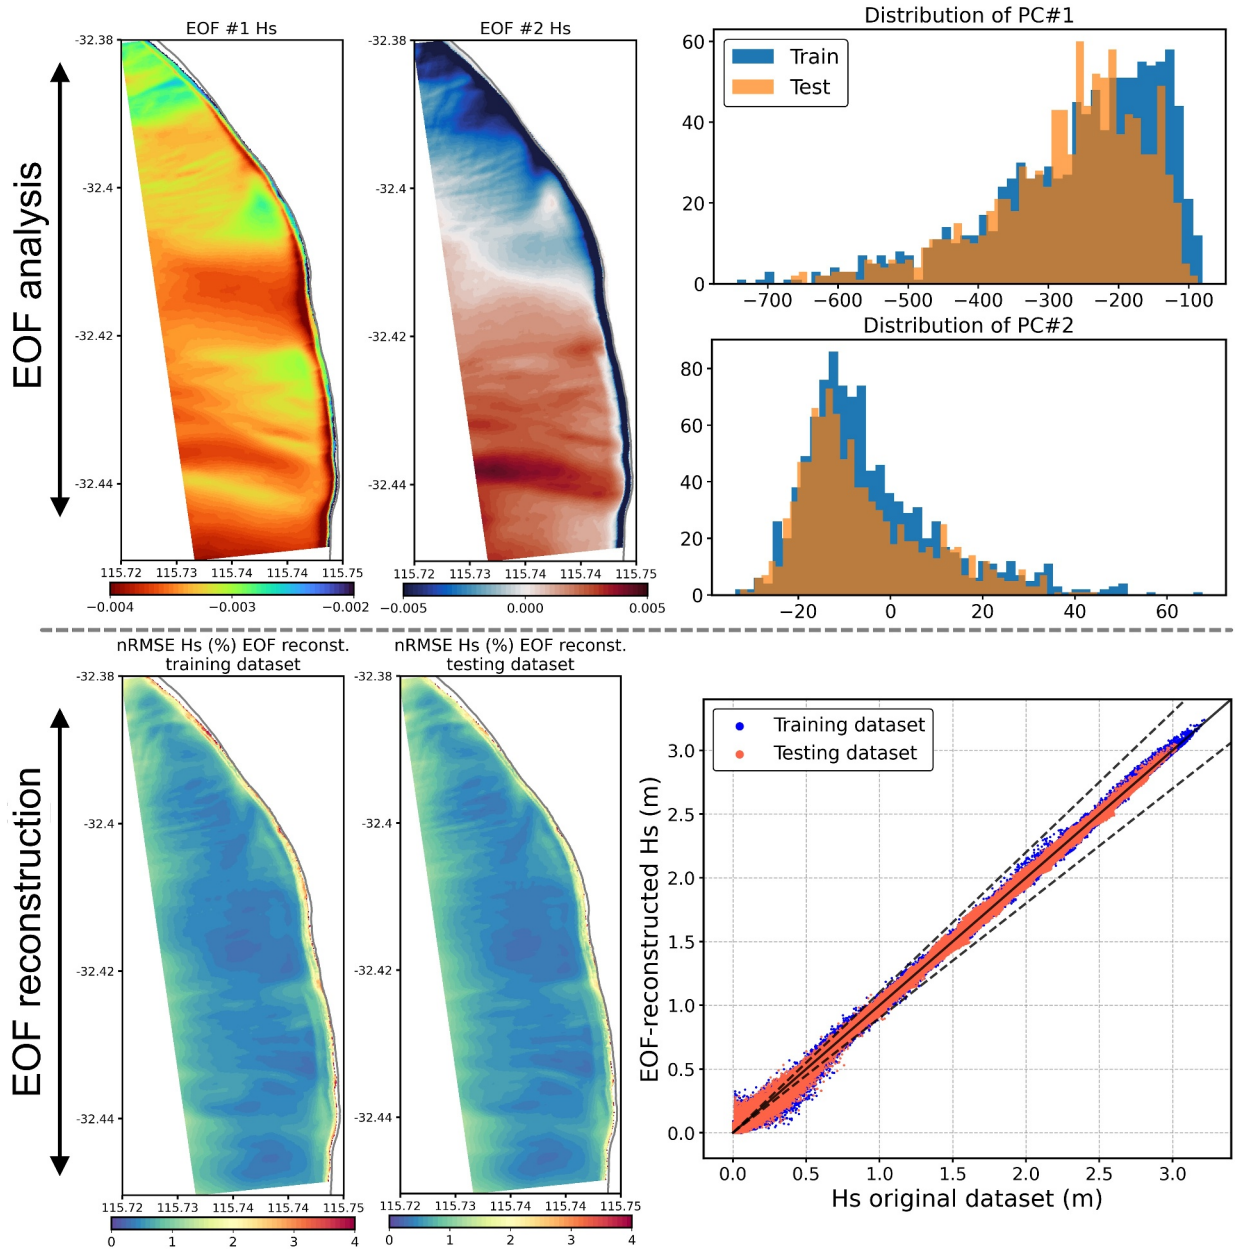

Figure 4. EOF analysis of H_s fields from XBeach-SB outputs. Top left: First two dominant EOF patterns, computed with the training data set. Top right: Histograms of the first two PCs for the training and testing data sets. Bottom left: Root-mean-square error normalized by the mean of the true H_s fields reconstructed with 20 EOFs. Bottom right: Scatter plot of all H_s values across all grid points and MDA cases.

data into these EOF patterns. Through this approach, the problem becomes building a ML model that can learn and predict the values of the PCs given the input parameters used to force the XBeach-SB simulations (wave boundary conditions and water levels). With the predicted PCs values, the original data can then be reconstructed using Equation 1.

As a demonstration of the methodology, the results for the short-wave significant wave height (H_s) and infra-gravity RMS (root mean square) wave height (H_{ig}) outputs will be presented hereafter. The results for other output variables can be found in Supporting Information S1. The first (dominant) two EOF patterns of the H_s output and their PCs distribution are presented in Figure 4. In addition, the original data was reconstructed using 20 EOFs, which ensures an accurate representation of the original data set (variance explained >99.9%). Increasing the number of EOFs did not result in a significant improvement in accuracy (see Figure S1 in Supporting

Information S1). The testing data can be reconstructed using the EOFs and PCs (see Equation 1), and the normalized root-mean-square error of the reconstructed H_s data set attains values of less than 3% overall, as shown in Figure 4, indicating the precision of the EOF-based decomposition.

As mentioned, the challenge becomes developing a ML model that uses the input parameters (namely wave boundary conditions and water levels) and predicts the PCs (20 in the case of H_s). Among the plethora of ML approaches available for such tasks, we decided to implement and compare three different approaches: (a) a deep neural network (DNN), (b) a Radial Basis Function (RBF) interpolator, and (c) a multi-variate linear regression (LR) model. The decision to use a DNN was guided by the seminal work of Hornik et al. (1989). Their study demonstrated that multi-layer feedforward neural networks are "universal approximators" of any measurable function, provided that enough hidden layers are selected. The authors argue that the ineffectiveness of a given DNN typically stems from three core issues: inadequate learning (which may be due to insufficient data or parameter fitting), a shortage of hidden layers within the network, or an inherent lack of relationship or connection between the inputs and outputs. Given these insights and considering the diverse options of neural networks potentially suitable for this application, we opted to employ a densely connected deep neural network. This choice is rooted in the principles laid out by Hornik et al. (1989), ensuring our model will be able to predict the PCs from the given input parameters. The specific details of the implementation of the neural network can be found in Appendix A. On the other hand, the RBF interpolator has also been successfully implemented in surrogate modeling studies in the past (e.g., Zornoza-Aguado et al., 2024; Ricondo et al., 2024, among many others). The RBF Interpolator employs a series of radially symmetric functions (in this case we use a Gaussian kernel) to represent a target function as a weighted sum of RBFs. These weights are adjusted during the training phase, allowing the model to effectively learn from the input forcing parameters and accurately predict the PCs. The choice of a Gaussian kernel was motivated by previous successful implementations of the RBF interpolator to predict nearshore oceanographic parameters (e.g., Guanche et al. (2013); Camus et al. (2013)). Lastly, the multivariate LR will serve as a fundamental baseline to discern the linearity in the relationships between inputs and outputs. These results will help in distinguishing scenarios where a linear model suffices from those requiring more complex, non-linear ML algorithms to accurately capture the underlying dynamics.

4. Results

4.1. Secret Harbour Domain

As outlined above, the DNN is trained to predict the 20 PCs of XBeach-SB's significant wave height (H_s) output using scaled input parameters (from the seven offshore boundary locations) on a data set comprising 1,000 MDA cases. It is then validated against a distinct test set of an equal number of MDA cases (see Figure 3). In this supervised learning approach, the neural network iteratively adjusts its weights and biases to minimize the loss function, using Adam (Kingma & Ba, 2014) as the optimization strategy to converge on a solution. Additionally, to prevent overfitting to the training data, the validation loss is monitored during the training phase, using the test data set. If the validation loss does not improve after 100 epochs, the training of the model stops.

After about 1,200 epochs, the model converges to a minimal loss in both the training and testing data sets, indicating a successful training of the network (see Appendix A). The DNN's performance is first evaluated by comparing the predicted PCs with the original ones from the test data set. This comparison is visually represented in Figure 5, which shows the scatter plots for the first nine original versus predicted H_s PCs. These plots underscore the DNN's proficiency in predicting the PCs, particularly for the first EOF modes. The first PC, being the dominant mode that accounts for the highest variance in the data set, shows points closely aligning with the 1:1 line, whereas for higher-order PCs increased scatter is observed, yet the comparison shows good agreement.

We repeat this process using the RBF interpolator as the ML model. The accuracy of the predictions of H_s PCs by the RBF interpolator is almost indistinguishable to that attained with the DNN, therefore these results are not shown. The final step involves using the predicted PCs to reconstruct the original test data set using Equation [1]. The reconstructed data set is then compared with the original XBeach-SB output for the test set.

Finally, we developed a LR model that links the input parameters with the PCs through a linear combination of the input parameters, each weighted by a coefficient, plus an intercept. A linear regression model was built for each PCs.

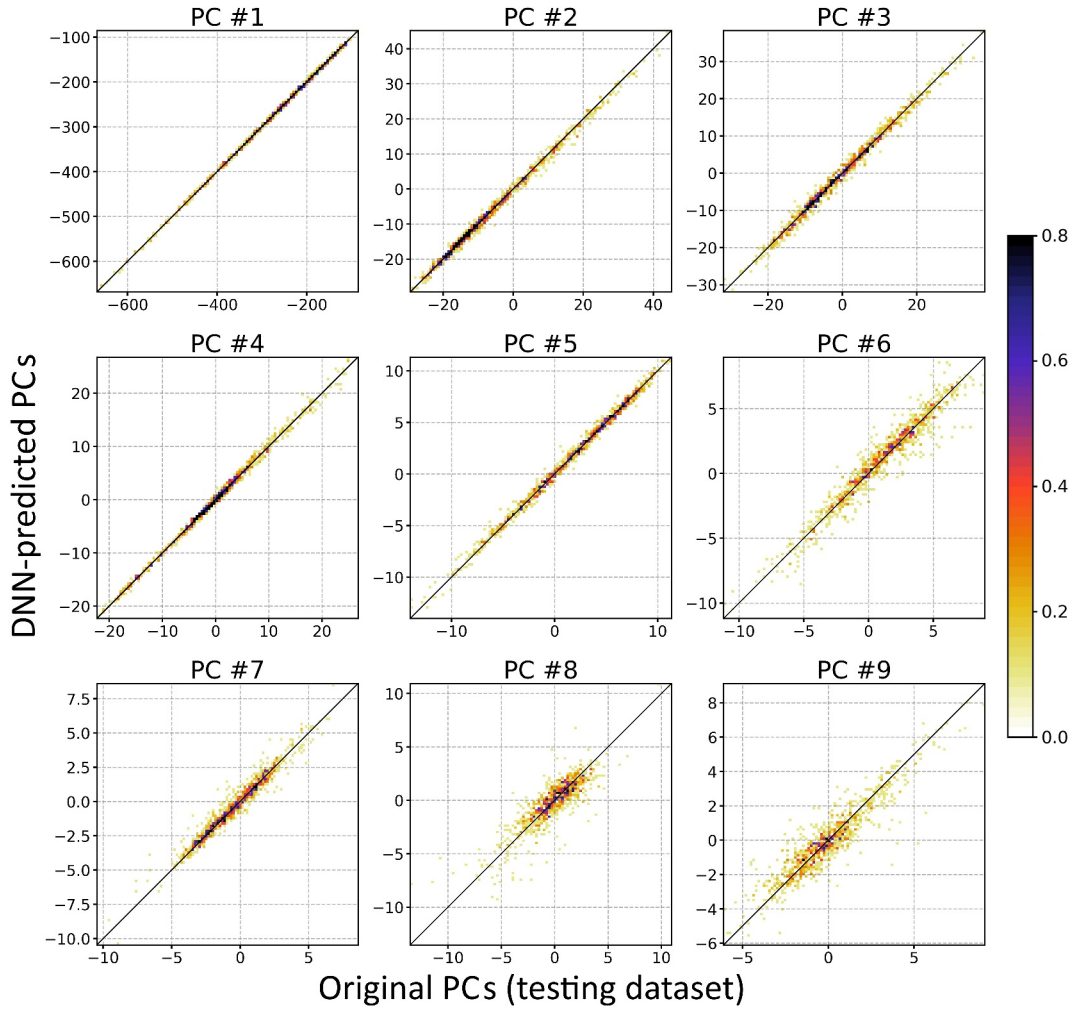

Figure 5. Scatter plots of the first 9 H_s PCs of the testing data set versus the PCs predicted by the neural network. The black line in each plot is the 1:1 line, and the colors indicate the density of points.

Figure 6 illustrates the accuracy of the reconstructed H_s: On the left, spatial maps show the normalized root-mean-squared error (nRMSE, normalized by the mean of the true values at each grid point) between the predicted wave height and that of the test data set, as rendered by the LR model (left), the RBF interpolator (middle) and the DNN (right). The LR model presents the highest errors, being around 3% in deep waters, increasing to well over 5% in shallower waters, particularly the surf-zone. Notably, for the RBF interpolator and the DNN, the reconstructed H_s shows errors below 2% in the deeper areas of the domain, increasing to approximately 3% in some places near the coastline. This indicates a high level of accuracy of the surrogate model, with the DNN being only marginally more accurate than the RBF interpolator, evidenced by the slightly lower nRMSE in the reconstructed H_s, particularly close to the coast. The right panel of Figure 6 features a scatter plot comparing actual versus reconstructed H_s across all grid points for all MDA cases in the test data set. There is a considerable scatter in the original versus reconstructed H_s for the LR model, whereas for the RBF interpolator and the DNN there is a closer correlation, with most points lying within ±10% of the original values, particularly for higher wave heights. However, discrepancies are higher for very low wave heights (below 0.5 m).

A variable of interest for many XBeach-SB applications is the infragravity wave height (H_{ig}). Given the complex and non-linear relationship between this variable, bathymetry and the forcing conditions such as integrated wave parameters and water levels, H_{ig} serves as an excellent candidate for assessing the predictive capabilities of our surrogate model. We applied the same methodology to the H_{ig} outputs as used previously for H_s: EOF analysis followed by the training of ML models (LR, RBF and DNN) to predict the PCs. For this variable, we increased the

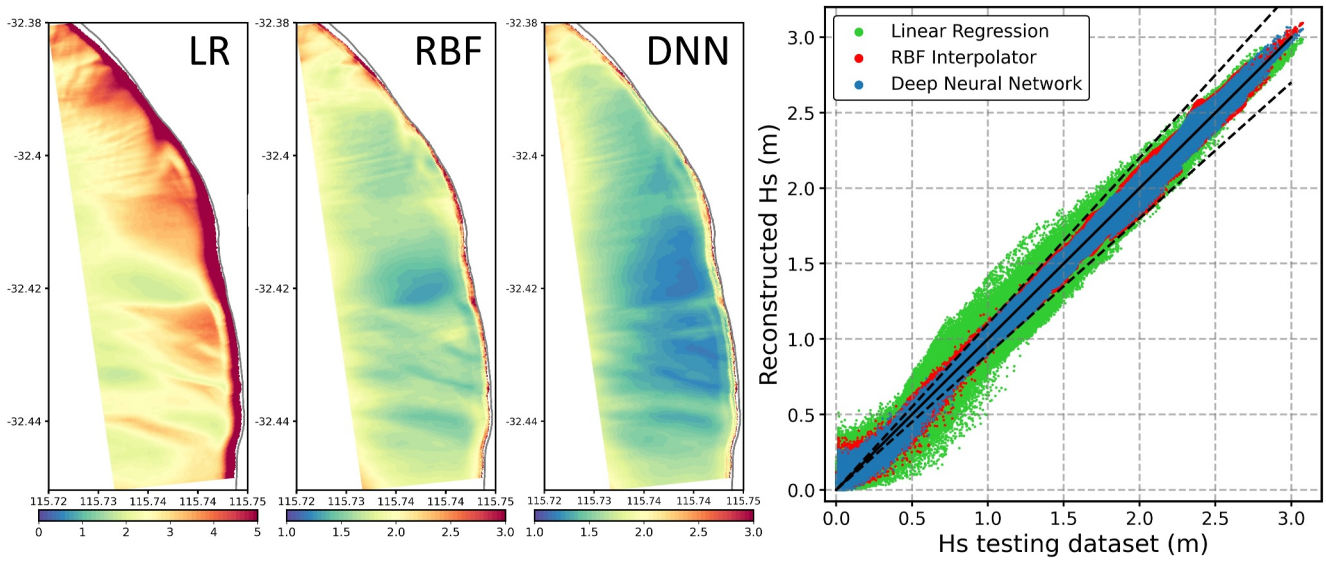

Figure 6. The spatial maps show the normalized root-mean-squared error (nRMSE) between the original XBeach-SB H_s output for the test data set and the reconstructed H_s (left: by the multi-variate linear regression model (LR); middle: by the RBF interpolator (RBF); right: by the deep neural network (DNN)). Notice the different ranges in the colorbars. The scatter plot shows the original versus reconstructed H_s output across all grid points for all MDA cases in the test data set. The black line indicates a 1:1 agreement, and the dashed lines represent a ±10% deviation from the 1:1 line.

number of EOFs to 30 to enhance reconstruction accuracy (see Figure S2 in Supporting Information S1). However, although increasing the number of EOFs increased the accuracy of the reconstruction, the precision of the surrogate model diminishes since the ML model makes inaccurate predictions of higher-order PCs, leading to increased errors in the reconstruction process. The first two EOF patterns and corresponding PCs for H_{ig}, along with the accuracy of the reconstruction using 30 EOFs, are detailed in Figures S3 and S4 of Supporting Information S1. Here we focus on the results of the final reconstruction with the surrogate model.

Figure 7 showcases the precision of H_{ig} reconstruction via the surrogate model. For the DNN-based reconstruction, the spatial distribution of nRMSE indicates a pattern akin to that observed with significant wave height

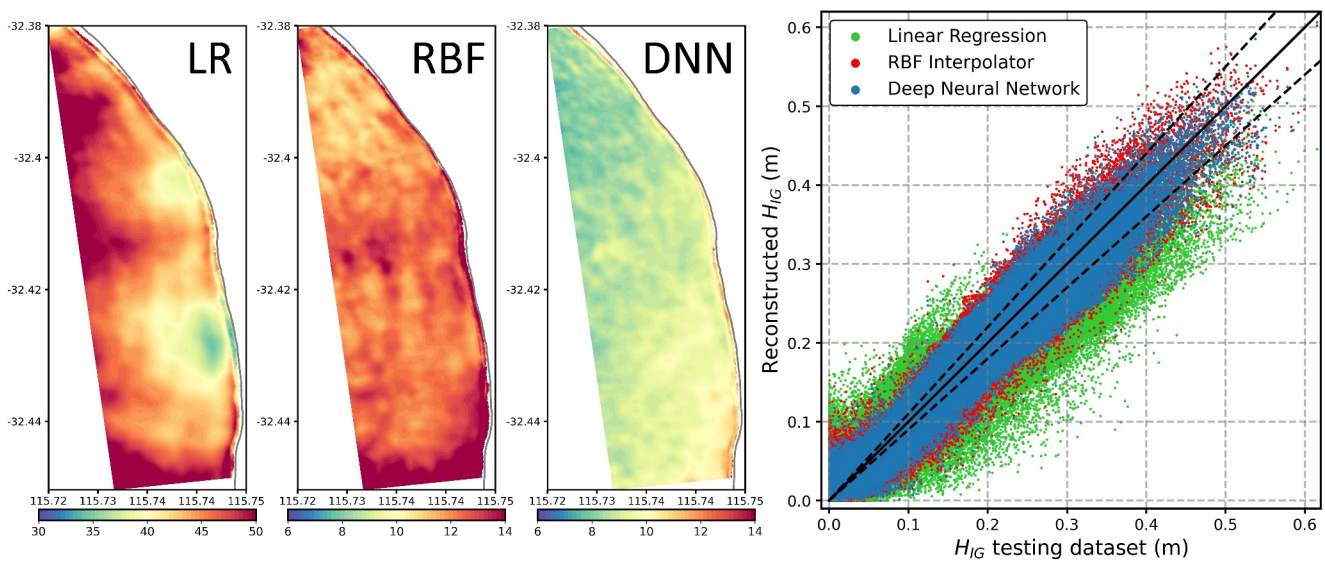

Figure 7. The spatial maps show the normalized root-mean-squared error (nRMSE) between the original XBeach-SB H_{ig} output for the test data set and the reconstructed H_{ig} (left: by the multi-variate linear regression model (LR); middle: by the RBF interpolator (RBF); right: by the deep neural network (DNN)). Notice the different ranges in the colorbars. The scatter plot shows the original versus reconstructed H_{ig} output across all grid points for all MDA cases in the test data set. The black line indicates a 1:1 agreement, and the dashed lines represent a ±10% deviation from the 1:1 line.

Table 1
Correlation Coefficient (r) and Root-Mean-Square Error (RMSE) Between Observed and Predicted Significant Wave Height (H_s), Infragravity Wave Height (H_{ig}), Mean Wavelength (L_{mean}), Fraction of Breaking Waves (Qb_{max}), Zonal and Meridional Components of Depth-Averaged Currents (ue_{mean} and ve_{mean}) for the Testing Data Set at Secret Harbour

	LR		RBF		DNN	
	r	RMSE	*r*	RMSE	*r*	RMSE
H_s	0.997	0.031 m	0.999	0.016 m	0.999	0.015 m
H_{ig}	0.963	0.013 m	0.991	0.006 m	0.992	0.005 m
L_{mean}	0.999	1.03 m	0.999	0.360 m	0.999	0.340 m
Qb_{max}	0.933	0.159	0.978	0.088	0.992	0.052
ue_{mean}	0.866	0.022 m/s	0.966	0.010 m/s	0.975	0.009 m/s
ve_{mean}	0.922	0.044 m/s	0.987	0.017 m/s	0.993	0.012 m/s

Note. Results for the multi-variate linear regression model (LR), the Radial Basis Functions interpolator (RBF) and the Deep Neural Network (DNN).

(H_s): lower errors (~8%) occur in the deeper regions of the domain, while errors increase nearer to the coast, reaching maximum nRMSE values of approximately 11%. On the other hand, the RBF model's reconstruction shows slightly higher values of nRMSE, generally ranging between 11% and 13% across the domain. The LR model, however, falls short in accurately predicting H_{ig}, achieving a nRMSE as high as 50% in certain domain areas, with the lowest observed nRMSE exceeding 30%. The scatter plot in Figure 7, encompassing all H_{ig} values from the testing data set, depicts a wider spread around the 1:1 line compared to the H_s results, with many grid points at certain time steps deviating beyond the 10% accuracy envelope. Despite the presence of this greater scatter, the absence of extreme outliers suggests that the models achieve some level of prediction skill for H_{ig}. Both the RBF and the DNN present an overall similar scatter, with the RBF having a slightly greater spread of data points outside the 1:1 line. The scatter points from the LR model's predictions underscore its relatively poor performance and unsuitability for this complex, non-linear variable.

Regarding the number of MDA cases selected for building the surrogate model, Figure S5 in Supporting Information S1 indicates that 1,000 cases is suitable to achieve a highly accurate surrogate model, and that increasing the number of MDA cases only brings marginal changes to the overall accuracy of the surrogate model. This means the selected 1,000 cases are representative of most conditions found in the historical record at this location, thus allowing the ML model to learn the full range of combinations necessary for accurate predictions. Beyond this point, adding more data becomes redundant, as it does not provide additional information that significantly improves model performance. This is consistent with previous studies using the MDA technique, such as Zornoza-Aguado et al. (2024) with 500 MDA cases, Ricondo et al. (2024) with 750 MDA cases, and Ricondo et al. (2023) with 800 MDA cases. Future work could develop location-dependent optimisation of the number of MDA cases, depending on the variability of local wave and sea level climate(s). Figures S6–S9 in Supporting Information S1 demonstrate the surrogate model's accuracy in predicting additional XBeach-SB wave variables, specifically mean wavelength (L_{mean}, Figure S6 in Supporting Information S1), fraction of breaking waves (Qb_{max}, Figure S7 in Supporting Information S1), and zonal and meridional components of depth-averaged Eulerian currents (ue_{mean} and ve_{mean}, Figures S8 and S9 in Supporting Information S1 respectively). While all ML models (LR, RBF and DNN) can predict the mean wavelength with a high level of accuracy, they perform less effectively in predicting the fraction of breaking waves. This variable is highly non-linear and discontinuous, which represents a significant challenge. Additionally, the fraction of breaking waves is dependent on the water level (wl) of each specific MDA case. The inter-tidal area, which is intermittently wet and dry, shows the highest values of Qb_{max} when the grid points in the surf zone are "wet", and no value is computed when they are dry. Thus, although wl is one of the training variables, the EOF decomposition is fundamentally limited by its inability to represent the wetting and drying processes in the model domain, and the related prediction of patterns of variability in the surf zone are less reliable. Finally, Figures S8 and S9 in Supporting Information S1 showcase the surrogate model's performance in predicting the zonal and meridional components of depth-averaged currents. The LR model again fails to accurately capture this variable, while the RBF and DNN models attain overall RMSE values below 0.04 m/s. The greatest discrepancies occur at lower current speeds, whereas higher speeds are predicted with greater accuracy. Table 1 summarizes the performance of the surrogate model for the Secret Harbour domain, presenting the correlation coefficients and the RMSE values of predicted versus actual XBeach-SB output variables in the testing data set.

The statistics presented above (as well as in Section 4.2) allow for an assessment of how well the surrogate model methods presented are able emulate the full numerical (XBeach-SB) model under the quasi-stationary boundary conditions associated with the MDA testing data set conditions. They do not allow for assessment of the surrogate model's ability to simulate realistic time-varying conditions, nor do they contextualize the surrogate model's errors relative to the XBeach-SB's predictive skill (i.e., validation against observations). The observations summarized in Section 2.1 (with locations show in Figure 1) are therefore compared to both XBeach-SB and surrogate model predictions with continuous time-varying boundary conditions. These conditions are derived from the same sources as the for the MDA simulations, but the simulations run continuously for a 2-week period

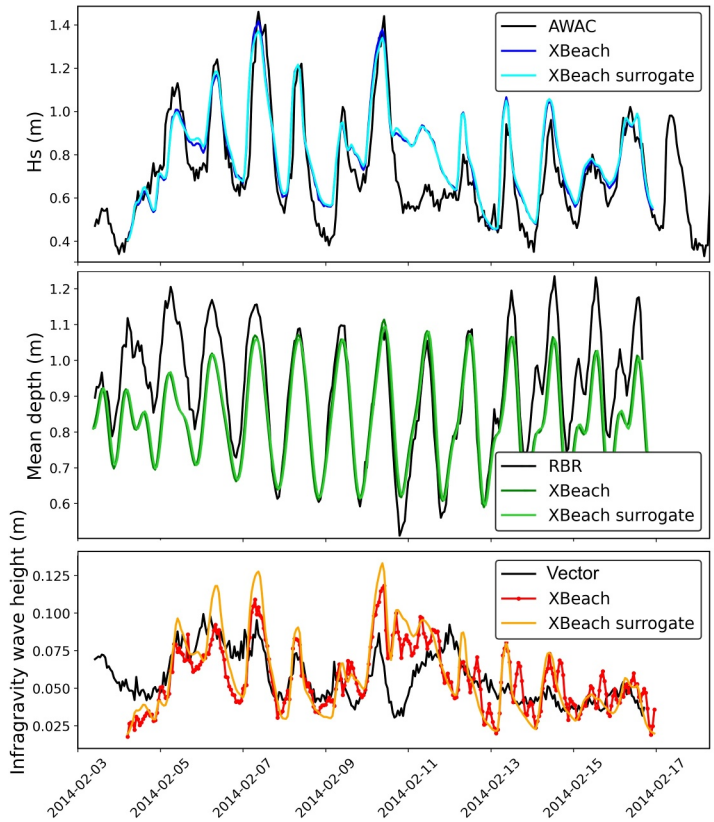

	r	RMSE
XBeach	0.881	0.131 m
XBeach surrogate	0.876	0.133 m

	r	RMSE
XBeach	0.803	0.145 m
XBeach surrogate	0.802	0.147 m

	r	RMSE
XBeach	0.469	0.013 m
XBeach surrogate	0.558	0.014 m

Figure 8. Verification of XBeach-SB and DNN surrogate model simulations for February 2014 at observations locations indicated in Figure 1: top) significant wave height (H_s) compared against AWAC measurements; middle) mean water depth compared against RBR measurements; bottom) infragravity wave heights (H_{ig}) compared against calculated H_{ig} from 1 Hz observations from a Vector Velocimeter. The tables on the right show the Pearson correlation coefficient (r) and the root-mean-square error (RMSE) values for the numerical model output (XBeach) and the surrogate model predictions (XB-surrogate).

in February 2014. Note that the input forcing for the surrogate model is the same as that for the numerical XBeach simulations.

The validation against the AWAC H_s observations shows that the XBeach-SB simulations capture the pattern of wave variability (Figure 8). There is a diurnal cycle in the H_s values that is related to variations in the wind associated with the local sea breeze (this variation is provided by the SWAN hindcast data used as wave boundary forcing). However, XBeach-SB tends to overestimate low H_s values, and while some peaks are accurately captured, many are slightly underestimated. Specifically, the largest discrepancies in H_s outputs from XBeach-SB simulations manifest around the 10th of February 2014, with a notable overestimation observed. These errors can be attributed primarily to inaccurate wave forcing from the offshore boundaries. The XBeach-SB simulation also reproduces the pattern of oscillation in the time-averaged surf zone water levels (wl) as represented by RBR; however, it underestimates them at the beginning and end of the observation period. Similarly, while this variable presents the highest levels of disagreement between the model and observations, this XBeach-SB configuration captures the magnitude of H_{ig} height (at the Vector location) and discerns most of its fluctuations, particularly during the first half of the simulated period.

Beyond errors attributable to boundary forcing, utilisation of JONSWAP spectra limits the ability of XBeach-SB to capture complex cases involving multi-modal wave spectra, or spectral shapes different than those of a JONSWAP spectrum. Additionally, nearshore morphological changes can have a significant impact on wave transformations; discrepancies between the bathymetric features represented by the LiDAR-derived survey data used and subsequent changes can also play a role in contributing to the observed errors. While there is undoubtedly room for improvement, for example, calibration of the bed friction coefficients and other physical

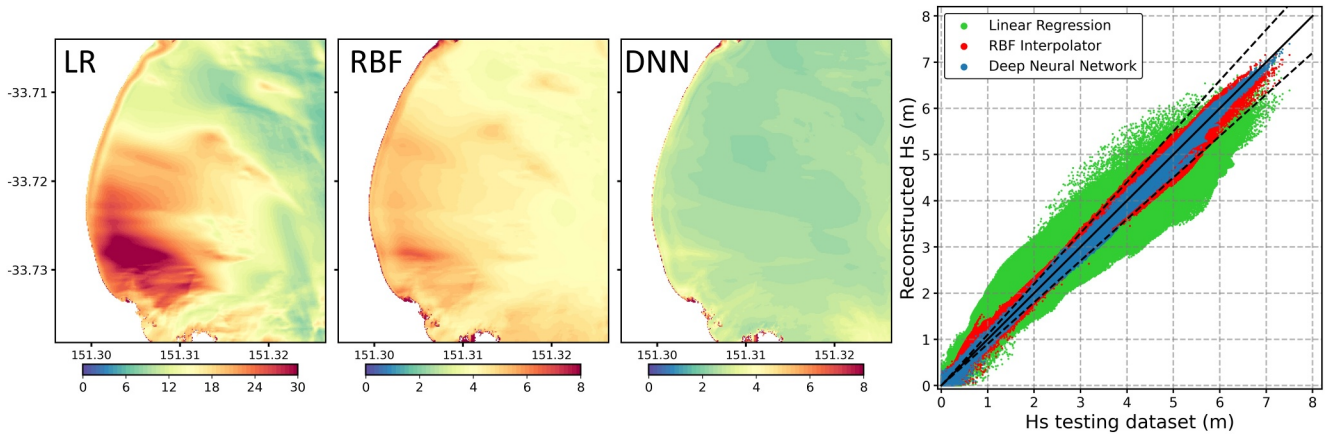

Figure 9. The three left maps show the normalized root-mean-squared error (nRMSE) between the original XBeach-SB H_s output for the test data set and the reconstructed H_s by the multi-variate linear regression model (LR), the RBF interpolator (RBF), and by the deep neural network (DNN). The scatter plot shows the original versus reconstructed H_s output across all grid points for all MDA cases in the test data set. The black line indicates a 1:1 agreement, and the dashed lines represent a ±10% deviation from the 1:1 line.

parameters (XBeach-SB implementations are known to benefit from extensive tuning, see for example Simmons et al., 2017), the aim of this study is not to carry out an exhaustive calibration/validation. Rather it is to compare the predictive skill of both XBeach-SB and develop the proposed surrogate modeling technique. Inspection of Figure 8 demonstrates that any differences in the predictive skill of the surrogate model are minimal compared to the skill of XBeach-SB to observations. The surrogate model closely replicates the XBeach-SB model output, with only minor discrepancies. The most significant differences are observed for the H_{ig}, which the surrogate model tends to overestimate. Although an extensive calibration of XBeach-SB is out of the scope of this paper, it is a fundamental step to ensure the accuracy of the surrogate model. Notably, ML methods can also be employed to perform calibrations of numerical models (see, for example, Itzkin et al., 2022; Simmons et al., 2017). Such methods could be applied prior to building the surrogate model as discussed in this paper. Alternatively, future work could involve incorporating the coefficient parameters as inputs to the surrogate model; this way, any sea state could be reconstructed under any combination of calibration coefficients, and through an optimisation algorithm, obtain the combination that yields the minimum errors when compared to measured data.

4.2. Narrabeen Domain

As described in Section 2, Secret Harbour and Narrabeen experience different tidal regimes and wave climates, so exactly the same methods described in the last section were applied to the Narrabeen XBeach-SB domain in order to evaluate the relocatability of the methods and test for sensitivities. That is, the same steps were taken: defining 1,000 training and 1,000 testing scenarios through MDA, conducting XBeach-SB simulations for each of these, carrying out an Empirical Orthogonal Functions analysis to selected gridded outputs, and training ML models to predict the PCs from the input forcing parameters.

Figure 9 demonstrates the accuracy of the reconstructed significant wave height (H_s) fields from the testing data set. As was the case for Secret Harbour, the LR model exhibits severe limitations, failing to accurately represent the wave heights inside the domain, with percentage errors ranging from 20% to 30% in shallow areas. On the other hand, both the DNN and the RBF interpolator show a similar performance, representing the wave heights inside the domain with a reasonable degree of accuracy, with maximum errors of around 8% in the surf zone. The DNN is marginally more accurate than the RBF for this domain.

Figure 10 shows the results of the surrogate model in representing infragravity wave height (H_{ig}), a variable characterized by its non-linear behavior compared to H_s. This non-linearity is evident in the performance of the linear regression model, which predicts H_{ig} with errors exceeding 60% in most of the domain, indicating the significant mismatch between predicted and actual values. Although the RBF model offers substantial improvements in accuracy with respect to the linear regression model, it still encounters errors of ∼25% in certain areas, highlighting its limitations in capturing the complex dynamics of infragravity waves. The DNN model

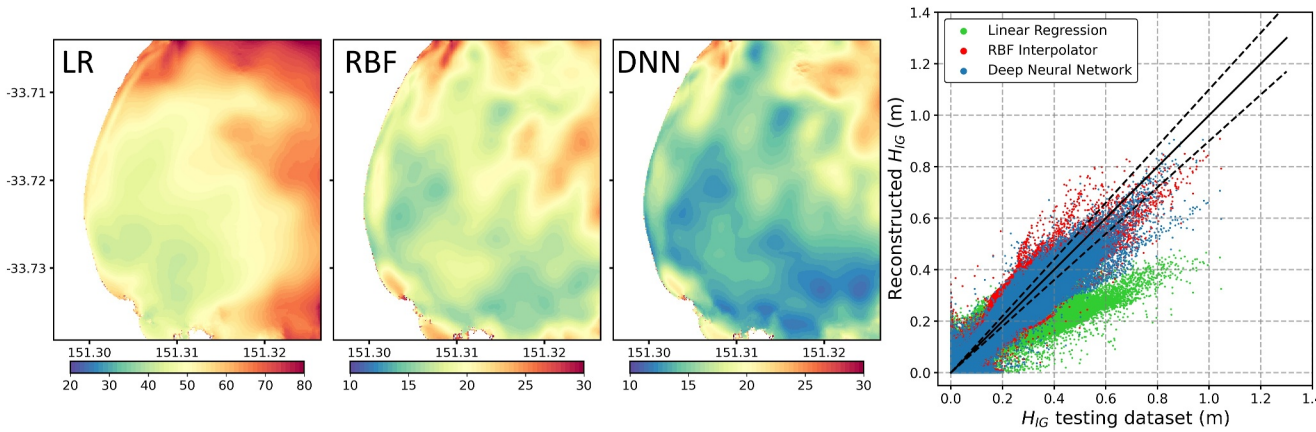

Figure 10. Same as Figure 9 but for infragravity wave height (H_{ig}).

exhibits the highest accuracy, outperforming the RBF model slightly, with an overall nRMSE ranging from 10% to 25% and displaying fewer outliers in the scatter plot.

5. Discussion

5.1. Surrogate Model Performance

This study presents a robust methodology for integrating high-fidelity numerical simulations with ML techniques to create a surrogate model capable of predicting key variables describing coastal hydrodynamics. This hybrid methodology leverages the XBeach-SB numerical model; when paired with efficient ML methods, we demonstrate the effectiveness of Deep Neural Networks (DNNs) and Radial Basis Functions (RBF) interpolators in capturing the complex coastal dynamics under varying conditions. Variables tested include significant wave height (H_s), infragravity wave height (H_{ig}), mean wavelength, fraction of breaking waves (Qb_{max}), and depth-integrated currents.

The DNN and RBF models exhibited high predictive accuracy for H_s, maintaining error margins well within ±10% across most scenarios. Specifically, maximum errors were around 3% at Secret Harbour and 8% at Narrabeen across the domains, and only a few grid points at certain times attaining errors of more than 10% (see scatter plots in Figures 6 and 9). The highest errors are found for H_s lower than 0.5 m, indicating areas where model performance could be further improved. In contrast, the multi-variate linear regression (LR) model presents significantly higher errors than the DNN or the RBF, particularly in shallow water areas, highlighting its inadequacy for non-linear processes such as wave breaking, currents and H_{ig} generation, which warrants the use of more complex ML models to represent the inherently complex coastal dynamics.

For H_{ig}, the DNN proved to be consistently slightly more accurate than the RBF interpolator. At Secret Harbour, it achieved errors generally around 8% with maximum nRMSE around 11% close to the coast. At Narrabeen the DNN's predictions of H_{ig} were less accurate, with errors ranging from 10% to 25% across the domain, however the scatter plot in Figure 10 indicates that the models are making reasonably accurate predictions of H_{ig}. In terms of mean wavelength, all ML models were able to make highly accurate predictions of it. On the other hand, the fraction of breaking waves (Qb_{max}) presents more challenges, due to the discontinuity in the grid points close to the coast, where Qb_{max} goes from one (on a "wet" grid point) to 0 (in a "dry" grid point corresponding to land). Despite this, our methodology can efficiently capture the main spatial patterns and intensity of Qb_{max}, with the highest errors being attained in the inter-tidal area where transitional water levels significantly influence the results.

The main strength of the hybrid surrogate model presented here is its computational efficiency, which enables extremely rapid prediction of gridded XBeach-SB output variables relative to the full dynamics numerical model. Such efficiency facilitates a range of applications, from enhancing nearshore forecasting capabilities (e.g., via ensembles), to generating extensive multi-year time series of XBeach-SB output fields, crucial for informed

Table 2
Estimated Run Times (Expressed as CPU Hours) for Creating a 40-Year Climatology Hourly Timeseries of XBeach Model Outputs, Using the Dynamical Numerical Model (Top Row) and the Hybrid Surrogate Model (Bottom Rows)

Dynamical numerical model	Running a full 40-year XBeach simulation	~1,600,000 CPU hr
Surrogate model	Carrying out 2,000 MDA XBeach simulations	~30,000 CPU hr
	Training the surrogate model	~0.15 CPU hr
	Creating 40 yearly NetCDF files of predicted outputs	~8 CPU hr

Note. This table highlights the computational efficiency of the hybrid approach relative to the fully numerical dynamical model.

decision-making in coastal management and planning. Table 2 compares the estimated computational time required to generate a 40-year deterministic simulation of XBeach outputs using both the fully numerical model and the surrogate model, which is a commonly used time frame to establish climatological baselines for extreme event analysis and coastal hazards assessment (e.g., Perez et al., 2017), except in certain areas such as those impacted by tropical cyclones (Haigh et al., 2014). Carrying out such long simulations with the fully dynamical numerical model can be quite computationally demanding, and the surrogate model can reduce this computational time by two orders of magnitude.

Moreover, the versatility of this methodology extends beyond the specific outputs discussed herein; it can predict additional variables and can be adapted to other numerical models, such as SWAN, WAVEWATCHIII, and SCHISM-WWM3 wave models, to develop corresponding surrogate models. This adaptability provides a robust framework for comprehensive coastal studies. Regarding the boundary conditions used to force the model, we utilized a 40-year hindcast of waves and water levels to create the MDA cases for training and testing, and (once the surrogate model has been built) to create a climatology of XBeach-SB outputs. As mentioned, alternative approaches include implementing the MDA on data sets generated through methods such as Latin Hypercube Sampling (e.g., Ricondo et al., 2024; Zornoza-Aguado et al., 2024), or applying the MDA to synthetic data sets created by stochastic climate emulators to generate plausible time series of future climate scenarios (Anderson et al., 2021; Wang et al., 2024). Additionally, the validation of the surrogate model can be made more robust in certain conditions by applying the MDA on the entire historical data set (if available) and implementing a K-fold cross-validation approach. Regardless of the method used, it is essential to ensure that the training and testing data sets encompass a wide range of potential conditions to enhance model generalizability.

5.2. Limitations of the Hybrid Surrogate Model

It is also important to highlight the limitations of this methodology for a proper interpretation of its outputs. First, the XBeach-SB simulations were carried out without the morphology option activated, resulting in a constant bathymetry across all simulations and surrogate model's predictions. This simplification does not reflect real-world conditions where morphological changes can significantly affect wave breaking patterns and wave propagation in shallow waters. In addition, this methodology does not contemplate changes in coastal structures, such as seawalls, breakwaters, and beach nourishment, which can also influence wave transformation, nearshore hydrodynamics and sediment transport. While the DNN can make accurate predictions of XBeach-SB outputs, the real-world accuracy of these predictions is inherently limited by the fidelity of the XBeach-SB configuration itself. Therefore, ensuring the accuracy of the XBeach-SB configuration is essential and might require a meticulous calibration prior to running the MDA cases to train and test the ML models.

The surrogate hybrid models presented here are predicated on training and testing using quasi-stationary simulations of XBeach-SB (i.e., the MDA training and test cases do not individually contain time-varying boundary conditions). Nevertheless, comparisons of surrogate model output generated using time-varying forcing conditions is in very good agreement with fully numerical XBeach-SB simulations using the same forcing (Figure 8), indicating the surrogate models' capability in this regard. The ability of the surrogate model to maintain accuracy using time-varying inputs, despite being trained with stationary inputs is likely due to the small domain size and quick "spin-up" period of the XBeach-SB model domains. In any case, the differences between surrogate and XBeach-SB time-varying predictions are significantly less than the difference between model predictions and in

situ observations. This indicates that (at least for this particular domain) the predictive skill of the surrogate models relative to XBeach-SB are considerably better than XBeach-SB's skill to observations.

In addition, the computational efficiency of this methodology is greatly facilitated by the EOF decomposition of the XBeach-SB gridded outputs. However, intermittent wetting and drying of grid points in the inter-tidal area presents challenges, as it introduces discontinuities that compromise the effectiveness of EOF analysis in accurately reducing and reconstructing the data. This limitation leads to errors in the EOF reconstruction of XBeach-SB output fields, which are then carried over to the ML models' predictions. In the two sandy beach domains studied, which have relatively small inter-tidal areas, it was feasible to identify and exclude these outliers. However, in more complex domains with extensive inter-tidal areas, this issue could significantly impact model accuracy. One potential solution is to bypass the EOF decomposition step and develop larger and more complex DNN architectures that are trained to predict the entire gridded output rather than just the PCs. While this approach might enhance prediction accuracy in the intertidal zone, it would tend to negate the computational efficiency provided by the EOF analysis. Specifically, the DNN built with the EOF approach contains approximately 400,000 trainable parameters, with training times ranging between 10 and 20 min (depending on the selected output variable) on a CPU-based environment. In contrast, constructing a DNN to predict the entire gridded output (without the EOF decomposition) involves around 30 million trainable parameters and can require 1–3 hr of training. While the latter approach can yield more accurate predictions in the surf zone, the observed improvements were overall marginal, leading us to conclude that the EOF-based approach offers a more effective balance between efficiency and accuracy for the two domains studied here.

6. Conclusions

The hybrid model developed in this study offers a computationally efficient and accurate approach for predicting gridded output fields from XBeach-SB simulations, which is valuable for a wide range of applications. The surrogate model facilitates the rapid downscaling of long-term nearshore wave conditions and nearshore hydrodynamics, crucial for effective coastal hazard analysis, management and planning, as well as enhancing coastal forecasting capabilities for operational and safety consideration, amongst other benefits.

Our results show that the DNN and RBF based models outperform traditional multi-variate linear regression, particularly for variables representing complex and non-linear coastal processes such as infragravity wave heights and waves breaking. However, recognizing the limitations of this approach (see the Discussion section) is essential for its effective application across varied coastal settings. Future research should focus on integrating morphological changes into the model to better mimic real-world coastal dynamics, and optimizing ML architectures to improve the surrogate model performance in complex coastal settings with significant inter-tidal variability. Additionally, expanding this hybrid technique to other coastal numerical models can broaden its utility, providing valuable tools for coastal engineers and scientists.

Appendix A: Neural Network Design

The neural network was constructed using the Keras python library. It is a feedforward neural network with multiple fully connected (dense) layers. Through iterative optimization and evaluation, the architecture that provided an optimum performance was: an input layer with as many neurons as input features (predictors), a Normalization layer, four hidden layers with 360 neurons, and an output layer matching the number of PCs to be predicted. The *leaky_relu* activation function was used for all dense layers except the output layer which had a linear activation function. For the compilation of the model, the Adam optimizer was used (Kingma & Ba, 2014), with the mean absolute error as the loss parameter and a learning rate of 1E−4. To enhance training efficiency, a callback was implemented to track the loss in the validation (testing) data set and stop the training of the model if the gains in accuracy in the validation set were not improved after 100 epochs. Finally, the input data was scaled using a *min-max* scaler. In cases where the range of the different PCs varies substantially, the PCs were standardized using the *StandardScaler* from the scikit-learn library; this ensures a balanced training of the network.

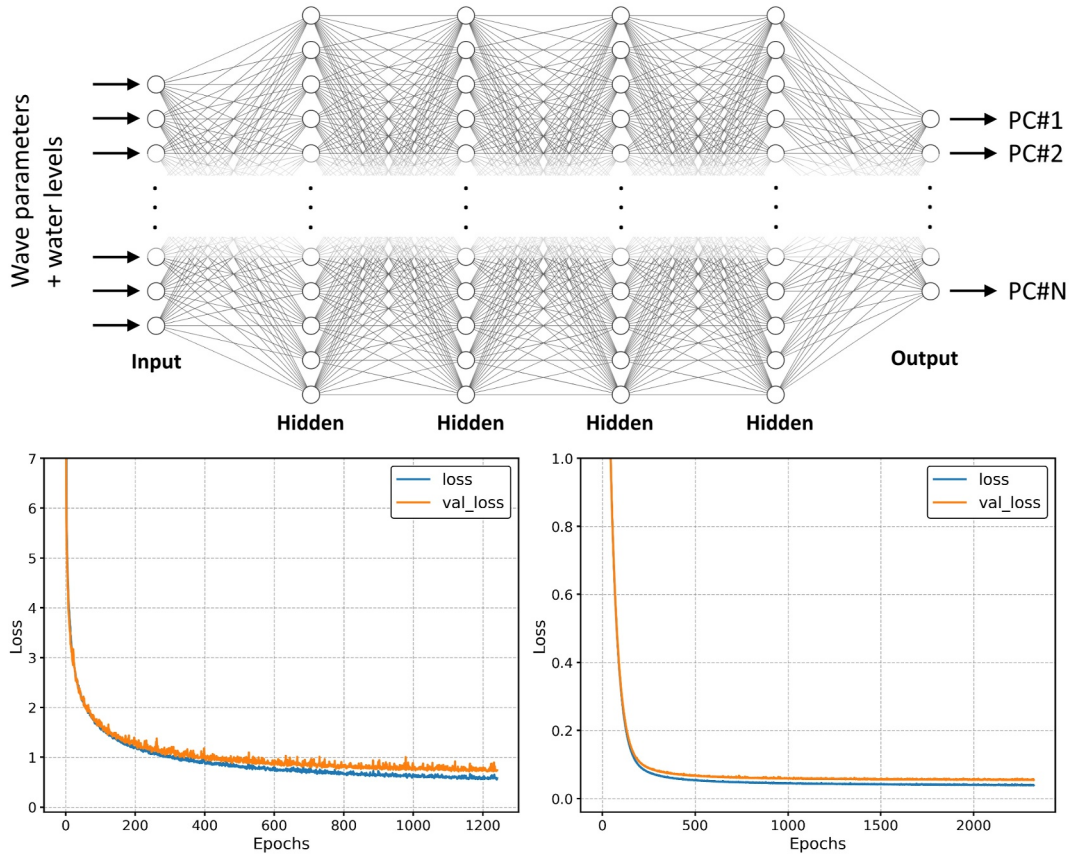

Figure A1. The top panel illustrates the architecture of the neural network employed in this study. The bottom panel depicts the evolution of the loss function over epochs during the neural network training, showcasing the learning progression (the significant wave height case on the left and the infragravity wave height case on the right).

Figure A1 shows a schematic of the network implemented in this study and the training-validation loss curves, illustrating the model's learning progress.

Data Availability Statement

The XBeach numerical model can be downloaded from https://download.deltares.nl/xbeach. The output from the 2,000 XBeach simulations carried out for the Secret Harbour domain (1,000 training and 1,000 testing simulations) can be obtained from Echevarria et al. (2025). The SWAN WAXA 500 m hindcast data used to create the wave boundary forcing at the Secret Harbour domain can be obtained from Trenham et al. (2024). For the Narrabeen domain, the CAWCR Wave Hindcast was used to generate the wave forcing files, and this can be downloaded from Durrant et al. (2019). The ORAS5 global ocean reanalysis used for sea level forcing in both simulations can be downloaded from Copernicus Climate Change Service, Climate Data Store, (2021). The tidal forcing for both simulations was generated using the "tpxo-tide-prediction" Python package (https://github.com/fwrnke/tpxo-tide-prediction), which utilizes the TPXO9-atlas models provided by Oregon State University. The Deep Neural Network was developed using the TensorFlow Python library (https://www.tensorflow.org/). The Maximum Dissimilarity Algorithm code and the Radial Basis Function code were sourced from the Bluemath initiative (https://github.com/GeoOcean/BlueMath_tk/tree/develop).

Acknowledgments
The authors wish to thank three anonymous reviewers for their insightful comments and constructive feedback. This research was supported by the Bluelink Partnership: a collaboration between the Australian Department of Defence, Bureau of Meteorology and CSIRO. This project was supported by resources and expertise provided by CSIRO IMT Scientific Computing. EE is supported by funding from an R + CSIRO Early Researcher Career Fellowship. LC acknowledges the funding from the Juan de la Cierva—Formacion FJC2021-046933-I/MCIN/AEI/10.13039/501100011033 and the European Union "NextGenerationEU"/PRTR.

References

Anderson, D. L., Ruggiero, P., Mendez, F. J., Barnard, P. L., Erikson, L. H., O'Neill, A. C., et al. (2021). Projecting climate dependent coastal flood risk with a hybrid statistical dynamical model. *Earth's Future*, *9*(12), e2021EF002285. https://doi.org/10.1029/2021EF002285

Björnsson, H., & Venegas, S. A. (1997). A manual for EOF and SVD analyses of climatic data. *CCGCR Report*, *97*(1), 112–134.

Camus, P., Mendez, F. J., Medina, R., & Cofiño, A. S. (2011). Analysis of clustering and selection algorithms for the study of multivariate wave climate. *Coastal Engineering*, *58*(6), 453–462. https://doi.org/10.1016/j.coastaleng.2011.02.003

Camus, P., Mendez, F. J., Medina, R., Tomas, A., & Izaguirre, C. (2013). High resolution Downscaled Ocean Waves (DOW) reanalysis in coastal areas. *Coastal Engineering*, *72*, 56–68. https://doi.org/10.1016/j.coastaleng.2012.09.002

Contardo, S., & Symonds, G. (2013). Infragravity response to variable wave forcing in the nearshore. *Journal of Geophysical Research: Oceans*, *118*(12), 7095–7106. https://doi.org/10.1002/2013JC009430

Contardo, S., Symonds, G., Segura, L. E., Lowe, R. J., & Hansen, J. E. (2019). Infragravity wave energy partitioning in the surf zone in response to wind-sea and swell forcing. *Journal of Marine Science and Engineering*, *7*(11), 383. https://doi.org/10.3390/jmse7110383

Copernicus Climate Change Service, Climate Data Store. (2021). ORAS5 global ocean reanalysis monthly data from 1958 to present. *Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*. https://doi.org/10.24381/cds.67e8eeb7. Accessed on 13-Aug-2024.

Durrant, T., Hemer, M., Smith, G., Trenham, C., & Greenslade, D. (2019). *CAWCR wave Hindcast—aggregated collection. v5*. CSIRO. Service Collection. Retrieved from http://hdl.handle.net/102.100.100/137152?index=1

Echevarria, E., Contardo, S., Perez Diaz, B., Hoeke, R., Leighton, B., Trenham, C., et al. (2025). Secret Harbour XBeach surrogate model training and testing simulations. *v1. CSIRO. Data Collection*. https://doi.org/10.25919/dpsr-rg31

Egbert, G. D., & Erofeeva, S. Y. (2002). Efficient inverse modeling of Barotropic Ocean tides. *Journal of Atmospheric and Oceanic Technology*, *19*(2), 183–204. https://doi.org/10.1175/1520-0426(2002)019<0183:EIMOBO>2.0.CO;2

Goldstein, E. B., Coco, G., & Plant, N. G. (2019). A review of machine learning applications to coastal sediment transport and morphodynamics. *Earth-Science Reviews*, *194*, 97–108. https://doi.org/10.1016/j.earscirev.2019.04.022

Guanche, Y., Camus, P., Guanche, R., Mendez, F. J., & Medina, R. (2013). A simplified method to downscale wave dynamics on vertical breakwaters. *Coastal Engineering*, *71*, 68–77. https://doi.org/10.1016/j.coastaleng.2012.08.001

Haigh, I. D., MacPherson, L. R., Mason, M. S., Wijeratne, E. M. S., Pattiaratchi, C. B., Crompton, R. P., & George, S. (2014). Estimating present day extreme water level exceedance probabilities around the coastline of Australia: Tropical cyclone-induced storm surges. *Climate Dynamics*, *42*(1–2), 139–157. https://doi.org/10.1007/s00382-012-1653-0

Hinkel, J., Feyen, L., Hemer, M., Le Cozannet, G., Lincke, D., Marcos, M., et al. (2021). Uncertainty and bias in global to regional scale assessments of current and future coastal flood risk. *Earth's Future*, *9*(7), e2020EF001882. https://doi.org/10.1029/2020EF001882

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, *2*(5), 359–366. https://doi.org/10.1016/0893-6080(89)90020-8

Itzkin, M., Moore, L. J., Ruggiero, P., Hovenga, P. A., & Hacker, S. D. (2022). Combining process-based and data-driven approaches to forecast beach and dune change. *Environmental Modelling & Software*, *153*, 105404. https://doi.org/10.1016/j.envsoft.2022.105404

Kim, T., & Lee, W. D. (2022). Review on applications of machine learning in coastal and ocean engineering. *Journal of Ocean Engineering and Technology*, *36*(3), 194–210. https://doi.org/10.26748/KSOE.2022.007

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. https://doi.org/10.48550/arXiv.1412.6980

McGranahan, G., Balk, D., & Anderson, B. (2007). The rising tide: Assessing the risks of climate change and human settlements in low elevation coastal zones. *Environment and Urbanization*, *19*(1), 17–37. https://doi.org/10.1177/0956247807076960

McKay, M. D., Beckman, R. J., & Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, *42*(1), 55–61. https://doi.org/10.1080/00401706.1979.1048975

Mortlock, T. R., & Goodwin, I. D. (2015). Directional wave climate and power variability along the Southeast Australian shelf. *Continental Shelf Research*, *98*, 36–53. https://doi.org/10.1016/j.csr.2015.02.007

Nieves, V., Radin, C., & Camps-Valls, G. (2021). Predicting regional coastal sea level changes with machine learning. *Scientific Reports*, *11*(1), 7650. https://doi.org/10.1038/s41598-021-87460-z

Perez, J., Menendez, M., & Losada, I. J. (2017). GOW2: A global wave hindcast for coastal applications. *Coastal Engineering*, *124*, 1–11. https://doi.org/10.1016/j.coastaleng.2017.03.005

Phillips, O. M. (1977). *The dynamics of the upper ocean* (Vol. 366). Cambridge University Press.

Pomeroy, A., Lowe, R., Symonds, G., Van Dongeren, A., & Moore, C. (2012). The dynamics of infragravity wave transformation over a fringing reef. *Journal of Geophysical Research*, *117*(C11). https://doi.org/10.1029/2012JC008310

Rautenbach, C., Trenham, C., Benn, D., Hoeke, R., & Bosserelle, C. (2022). Computing efficiency of XBeach hydro-and wave dynamics on Graphics Processing Units (GPUs). *Environmental Modelling & Software*, *157*, 105532. https://doi.org/10.1016/j.envsoft.2022.105532

Ricondo, A., Cagigal, L., Pérez-Díaz, B., & Méndez, F. J. (2024). HySwash: A hybrid model for nearshore wave processes. *Ocean Engineering*, *291*, 116419. https://doi.org/10.1016/j.oceaneng.2023.116419

Ricondo, A., Cagigal, L., Rueda, A., Hoeke, R., Storlazzi, C. D., & Méndez, F. J. (2023). HyWaves: Hybrid downscaling of multimodal wave spectra to nearshore areas. *Ocean Modelling*, *184*, 102210. https://doi.org/10.1016/j.ocemod.2023.102210

Roelvink, D., Reniers, A., Van Dongeren, A. P., De Vries, J. V. T., McCall, R., & Lescinski, J. (2009). Modelling storm impacts on beaches, dunes and barrier islands. *Coastal Engineering*, *56*(11–12), 1133–1152. https://doi.org/10.1016/j.coastaleng.2009.08.006

Segura, L. E., Hansen, J. E., Lowe, R. J., Symonds, G., & Contardo, S. (2018). Shoreline variability at a low-energy beach: Contributions of storms, Megacusps and sea-breeze cycles. *Marine Geology*, *400*, 94–106. https://doi.org/10.1016/j.margeo.2018.03.008

Serafin, K. A., Ruggiero, P., Parker, K., & Hill, D. F. (2019). What's streamflow got to do with it? A probabilistic simulation of the competing oceanographic and fluvial processes driving extreme along-river water levels. *Natural Hazards and Earth System Sciences*, *19*(7), 1415–1431. https://doi.org/10.5194/nhess-19-1415-2019

Simmons, J. A., Harley, M. D., Marshall, L. A., Turner, I. L., Splinter, K. D., & Cox, R. J. (2017). Calibrating and assessing uncertainty in coastal numerical models. *Coastal Engineering*, *125*, 28–41. https://doi.org/10.1016/j.coastaleng.2017.04.005

Smallegan, S. M., Irish, J. L., Van Dongeren, A. R., & Den Bieman, J. P. (2016). Morphological response of a sandy barrier island with a buried seawall during Hurricane Sandy. *Coastal Engineering*, *110*, 102–110. https://doi.org/10.1016/j.coastaleng.2016.01.005

Smith, G. A., Hemer, M., Greenslade, D., Trenham, C., Zieger, S., & Durrant, T. (2021). Global wave hindcast with Australian and Pacific Island Focus: From past to present. *Geoscience Data Journal*, *8*(1), 24–33. https://doi.org/10.1002/gdj3.104

State Government of NSW and NSW Department of Climate Change, Energy, the Environment and Water. (2023). NSW bathymetry sourced from Multibeam and marine Lidar surveys, accessed from the Sharing and Enabling Environmental Data Portal. Retrieved from https://datasets.seed.nsw.gov.au/dataset/aa8f268e-a23d-4d27-b046-f60c45f8349b. date accessed 2025-02-28.

Tausía, J., Delaux, S., Camus, P., Rueda, A., Méndez, F., Bryan, K. R., et al. (2023). Rapid response data-driven reconstructions for storm surge around New Zealand. *Applied Ocean Research*, *133*, 103496. https://doi.org/10.1016/j.apor.2023.103496

Trenham, C., Leighton, B., Seers, B., & Hoeke, R. (2024). SWAN WAXA 500m wave hindcast. v2. *CSIRO. Data Collection*. https://doi.org/10.25919/se4e-g288. Accessed on 13-Aug-2024.

Turner, I. L., Harley, M. D., Short, A. D., Simmons, J. A., Bracs, M. A., Phillips, M. S., & Splinter, K. D. (2016). A multi-decade dataset of monthly beach profile surveys and inshore wave forcing at Narrabeen, Australia. *Scientific Data*, *3*(1), 1–13. https://doi.org/10.1038/sdata.2016.24

Tu'uholoaki, M., Espejo, A., Wandres, M., Singh, A., Damlamian, H., & Begg, Z. (2023). Quantifying mechanisms responsible for extreme coastal water levels and flooding during severe tropical cyclone Harold in Tonga, Southwest Pacific. *Journal of Marine Science and Engineering*, *11*(6), 1217. https://doi.org/10.3390/jmse11061217

Wang, Z., Leung, M., Mukhopadhyay, S., Sunkara, S. V., Steinschneider, S., Herman, J., et al. (2024). A hybrid statistical–dynamical framework for compound coastal flooding analysis. *Environmental Research Letters*, *20*(1), 014005. https://doi.org/10.1088/1748-9326/ad96ce

Wang, Z., Leung, M., Mukhopadhyay, S., Sunkara, S. V., Steinschneider, S., Herman, J., et al. (2025). Compound coastal flooding in San Francisco Bay under climate change. *npj Natural Hazards*, *2*(1), 3. https://doi.org/10.1038/s44304-024-00057-0

Zornoza-Aguado, M., Pérez-Díaz, B., Cagigal, L., Castanedo, S., & Méndez, F. J. (2024). An efficient metamodel to downscale total water level in open beaches. *Estuarine, Coastal and Shelf Science*, *299*, 108705. https://doi.org/10.1016/j.ecss.2024.108705