Coastal Engineering 205 (2026) 104911

Contents lists available at ScienceDirect

Coastal Engineering

journal homepage: www.elsevier.com/locate/coastaleng

Beach monitoring: Do we need to survey to spring low tide?

Samuel Rose [a,*], Chris Blenkinsopp [a], Andrew Barnes [b], William Russell [c], Charlie Thompson [d,e]

[a] Department of Architecture & Civil Engineering, University of Bath, Claverton Down, Bath, BA2 7AY, Somerset, UK
[b] Department of Computer Science, University of Bath, Claverton Down, Bath, BA2 7AY, Somerset, UK
[c] Wales Coastal Monitoring Centre, The Alps, Wenvoe, CF5 6AA, Cardiff, UK
[d] Channel Coastal Observatory, National Oceanography Centre, European Way, Southampton, SO14 3ZH, Hampshire, UK
[e] School of Ocean and Earth Science, University of Southampton, European Way, Southampton, SO14 3ZH, Hampshire, UK

ARTICLE INFO

Keywords:
Neural networks
LSTM
MLP
CNN
Beach profiles
Beach surveys
Xbeach

ABSTRACT

When collecting coastal monitoring data, it is common practice to survey down to spring low tide to capture the maximum extent of the exposed subaerial beach. However, collecting topographic beach data is possible for only a few days per month. By reducing the seaward extent of the survey, the incurred costs and risks to the survey schedule could be greatly reduced. However, this would result in information loss at the lowest extremes of the subaerial beach. This study assesses the feasibility of predicting this part of the beach using deep learning neural networks based on partial beach profile data. A range of network architectures were tested alongside linear extrapolation, which was used as a baseline model. Each model was tested on three beaches with varying morphology, ranging from steep (reflective) to mildly sloping (dissipative). The presence of morphological features was found to play a dominant role in the accuracy of the predicted profiles; profiles with more pronounced cross-shore morphological features, such as sandbars, produced the highest error. While local connectivity of each network architecture was found to be the key factor in producing realistic profiles, the 1D Convolutional Neural Network was found to be the most effective with an average RMSE of between 0.026–0.119 m. This RMSE is not substantially larger than the vertical accuracy of current survey techniques (0.03 m), and the study found that errors of this magnitude have negligible effects when the survey data is used to calculate beach volumes and conduct numerical wave runup analysis to assess coastal flood risk.

1. Introduction

Coastal monitoring programmes are designed to obtain geomorphological and/or hydrodynamic information about our coastlines to help understand coastal change and provide the evidence needed to aid coastal management decision making. Traditionally, one of the key coastal monitoring datasets consists of subaerial beach topography (landward of the shoreline), and while complete digital elevation models captured using drone and plane mounted Lidar or photogrammetry are gaining popularity, some of the most useful datasets comprise discrete beach profiles captured on timescales of months/years at repeated locations.

The National Network of Coastal Monitoring Programmes and the Wales Coastal Monitoring Centre are responsible for collecting and compiling the majority of beach profile data in England and Wales. This work is funded by DEFRA via the Environment Agency and the Welsh Government (NNRCMP, 2023). Their topographic survey specifications require profiles to extend to the level of mean low water spring tides (MLWS). While this approach allows collection of the full cross-shore extent of the subaerial beach, it means that surveys can be carried out on only 2–3 days per month, making data collection logistically challenging, costly and vulnerable to interruption or cancellation due to adverse weather. Reducing the seaward extent of the survey specification would reduce such risks by increasing the number of potential survey windows and possibly enabling more frequent surveys at a larger number of locations. For example, relaxing the survey specifications to MLW would increase the proportion of potential survey days from 7% to 50%, greatly increasing the efficiency of data collection and potentially opening up the potential to increase the spatial or temporal resolution of data collection; however, this modification would result in the loss of data from the lower intertidal zone.

While the scientific community uses beach profile data for a wide range of purposes, the most common uses for the coastal managers who typically commission long-term monitoring programmes are: (1) to track beach volume over time to warn of concerning sediment loss and (2) to assess coastal flood risk due to wave runup and overtopping (Harley et al., 2011). The consequences of reducing the cross-shore

* Corresponding author.
E-mail address: str24@bath.ac.uk (S. Rose).

https://doi.org/10.1016/j.coastaleng.2025.104911
Received 1 August 2025; Received in revised form 25 October 2025; Accepted 12 November 2025
Available online 24 November 2025

extent of beach surveys on the utilisation of data have not previously been considered in detail. Saville, Jr. and Caldwell (1952) and Silveira et al. (2013) assessed the effects that alongshore profile spacing had on beach volume calculation; however, neither considered the effects of cross-shore profile extent during their study. In their study of the implications of changing morphology on wave runup, Phillips et al. (2017, 2020) found that wave runup predictions using XBeach were sensitive to upper beach morphology, but the impact of morphology in the lower intertidal zone was not explicitly addressed.

The concept of a beach equilibrium profile, around which the envelope of natural beach profiles fluctuates over time, is widely used within coastal engineering and dates back to Fenneman (1902). While beaches are constantly evolving in response to changing wave and tidal conditions, their overall form tends to be very consistent, typically following a monotonic concave-up mean profile as first quantified by Bruun (1954) who found that the average profile shape for beaches in California and Denmark followed a power-law. Morphological features such as sandbars represent perturbations from this overall form.

Given the relatively consistent mean shape of typical beach profiles, this study explores the potential to reduce the cross-shore extent of beach profile surveys to enhance the efficiency of coastal monitoring programmes. In order to mitigate the loss of data associated with reducing the profile extent, the potential to predict the lower intertidal portion of the beach profile using available 'partial' profile data, applying linear extrapolation and machine learning methods trained on historical beach profile data is explored.

Previous efforts to predict beach profiles using machine learning have largely focused on predicting seasonal beach morphology variation driven by changes in external forcing (waves, tides, wind, etc.) (e.g. Hashemi et al., 2010; Li et al., 2020; López et al., 2018a,b). All of these studies concluded that neural network models performed significantly better than comparable mathematical, numerical, and probabilistic methods. However, all investigated a single location with relatively limited profile variations.

This study evaluates the feasibility of predicting partial beach profiles using simple mathematical models and the use of more complex supervised deep learning models with a view to increasing the efficiency of beach monitoring. Neural networks have been used in other fields to predict partial data sets, such as future stock prices (Moghar and Hamiche, 2020), however, there has been no research on the application of neural networks to extrapolate and extend partial beach profiles, reducing the survey requirement. Here, we explore, for the first time, the ability of a variety of neural network architectures to predict lower intertidal beach profile data based primarily on a "current" partial beach profile and historic beach topography on a range of beach types, and quantify the impact on volumetric and overtopping calculations.

Section 2 provides details on the study location, data used, and predictive models tested. The results of the study are presented in Section 3. Finally, the implications of using the proposed machine learning approach for coastal engineers and managers are discussed in Section 4, and conclusions are drawn in Section 5.

2. Methodology

2.1. Study locations

Pre-existing topographic survey data collected by the Southwest and Southeast regional coastal monitoring programmes (Channel Coastal Observatory, 2021), were used for this study. Three UK beach locations with a comprehensive record of beach survey data, each with different morphological features and geographic locations. These beaches were Gwithian, Wittering, and Brighton, representing high-energy dissipative, moderate-energy dissipative and reflective beach types respectively.

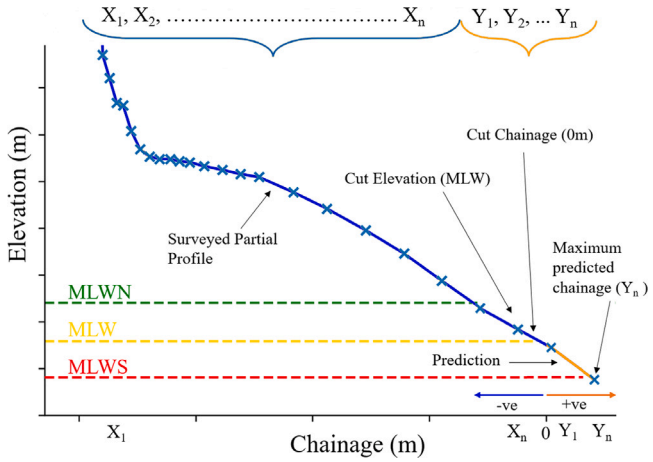

Fig. 1. An example beach profile. The blue line represents the surveyed partial profile, with measured points indicated by the crosses. The orange line represents the section of the beach profile that the study aims to predict. Negative chainages (X_1, X_2, …, X_n) are used to describe the measured region of the profile. Positive chainages (Y_1, Y_2, …, Y_n) represent the predicted region. The cut point that separates the two occurs at mean low water (MLW) at 0 m elevation and 0 m chainage.

A survey unit was chosen at each beach location, with each unit consisting of many transects. Two types of surveys were used in the study: baseline surveys, which captured data along all transects within the unit, and interim surveys, which recorded a smaller, representative subset. To maximise the training dataset, all surveys conducted prior to April 2023 were collated to develop a profile database, with many transects surveyed multiple times over the study period. While only a portion of this data could be used due to the constraints discussed in Section 2.2, the final dataset used to train and assess the models comprised 481 surveyed profiles at Gwithian (7a7A2-6/7), 268 at Wittering (5aSU02), and 155 at Brighton (4dSU13).

These profiles were surveyed using an RTK GPS, terrestrial laser scanner or UAV LiDAR, with vertical accuracy within 30 mm (NNR-CMP, 2023). Individual profiles were surveyed at varying intervals, determined by their significance and the type of survey commissioned, in accordance with a risk-based survey schedule.

2.2. Profile data

The goal of this study is to investigate the potential to predict the seaward part of a beach profile based on a partial survey down to a chosen tide level (hereafter termed "cut level") using a neural network or simple extrapolation. An example of a complete survey profile extending down to Mean Low Water Springs (MLWS) can be seen in Fig. 1. For this study, the cut elevation used to delineate the partial surveyed profile and the unsurveyed lower extremes of the subaerial beach was chosen as Mean Low Water (MLW) to balance additional survey time with the quantity of survey data lost. This choice allows surveys to be completed at low tide on approximately 50% of days, greatly increasing the available survey time compared to the current MLWS requirement (7% of days).

All survey elevations were referenced to Mean Low Water (MLW), defined as 0 m, with elevations representing vertical distance above or below this level. The chainage along each profile was transformed such that the origin (0 m) lies at the intersection between each profile and the cut elevation (MLW). Negative chainages (X_1, X_2, …, X_n) represent the surveyed partial profile and positive chainages (Y_1, Y_2, …, Y_n) represent the prediction region. In addition, since the distance between collected survey points varies spatially across the beach, with point density decreasing seaward, all profiles were interpolated onto

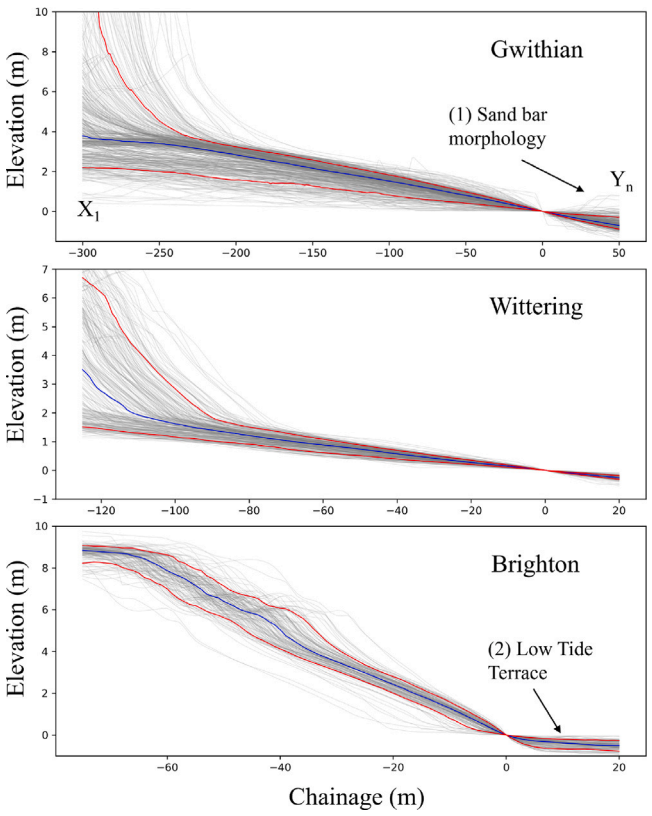

Fig. 2. All profiles used in the study at all three locations. The data has been preprocessed and discretised, with all profiles transformed to cross the origin at MLW. (Red) Shows the 95th and 5th percentiles of the profiles at each location. (Blue) Shows the mean profile at each location. X_1 and Y_n mark the start and the end of the profiles respectively.

a 1 m 1D grid. This preprocessing was completed to eliminate the potential for model bias, where a model's success would be driven by the accuracy of predictions in areas of higher survey density at the possible expense of less densely surveyed regions of the profile. A 1 m grid was chosen as a balance between minimising profile smoothing during preprocessing and managing the computational load during network training, ensuring the scalability of the method.

Due to the fixed shape of most neural networks, some profiles had to be trimmed or removed to be included in the training data for each model. Measured profiles that did not extend to the maximum predicted chainage (Y_n in Figs. 1 and 2) could not be used in training so were discarded in the prepossessing stage. Additionally, profiles that extended further than Y_n or began before the starting chainage (X_1 in Figs. 1 and 2) had to be shortened. The maximum chainage (Y_n) and starting chainage (X_1) were chosen at each location to find the optimal balance between the quantity/length of training data and the length of the profile predicted. Two additional characteristics of each profile were extracted from the survey data alongside the topographic information. A profile ID was assigned to each survey profile location. Profiles are surveyed in the same position periodically to allow temporal patterns to be extracted from sequential surveys; inclusion of this parameter aids the models to learn about local trends within the chosen beaches. Each profile location was given a unique number, which was later normalised for heightened network performance. Secondly, each survey was given a normalised seasonal value ($\hat{S}_n$) depending on the date that the survey was undertaken. Beach profiles typically exhibit seasonal morphological behaviour, and so this value provides a means to capture this within the model. ($\hat{S}_n$) was calculated using Eq. (1), where D_O is the ordinal date. $\hat{S}_n$ is a measure of how close the survey date is to the winter equinox, thus values close to 0 represent surveys conducted near the summer equinox, and values close to 1 represent surveys conducted near the winter equinox.

$$\hat{S}_n = \frac{|D_O - 183|}{183} \qquad 0 \leq \hat{S}_n \leq 1 \tag{1}$$

Beach profiles exhibit a spatially sequential structure, with cross-shore elevations displaying coherent morphological features such as berms, bars, and troughs. Morphological characteristics within a profile are often interrelated, with features often co-varying (e.g. steeper foreshore slopes are frequently associated with narrower berms or deeper troughs). These spatially consistent patterns reflect the primary controls on beach morphology, with hydrodynamic forces such as waves and tides, together with sediment composition, driving sediment transport and shaping morphological change. Consequently, these morphological features are often consistent across many profiles at the same study location. These relationships have been extensively studied and documented in a large body of literature that have focused on the morphological drivers, relationships and, evolution of beach profiles (Dean, 1977; Wright and Short, 1984; Masselink and Short, 1993; Scott et al., 2011; Castelle and Masselink, 2023). Building on this established understanding of profile morphology, this study examines both dissipative and reflective beaches allowing the exploration a wide spectrum of features and profile shapes.

All profiles have been plotted in Fig. 2 to highlight the morphological variability within and across study locations. Gwithian exhibits longshore bar systems with pronounced crests and troughs (1), producing highly variable profiles below MLW. Wittering has a more stable morphology with less pronounced bars and an intermittent low tide terrace. Brighton is characterised by an almost ubiquitous low tide terrace (2), yielding relatively uniform slopes and minimal variation along the profiles.

2.3. Profile prediction model architecture

To determine the ability of a model to predict the seaward extent of a beach profile based on partial beach profile information, five different modelling approaches were tested. Simple linear extrapolation was employed as a baseline alongside four largely different neural network architectures. In this study, LSTM, CNN, and MLP models were employed as representative architectures for sequential, spatial, and fully connected learning, respectively. These widely used neural networks were chosen to enable a comprehensive evaluation and comparison of their predictive performance.

2.3.1. Linear extrapolation (LE)

Linear extrapolation was selected as the baseline model against which to assess the performance of the neural networks. The empirical models proposed by Bruun (1954), Dean (1977, 1991), and Vellinga (1982) tend towards linear at the seaward extreme of the subaerial beach profile, thus, a linear extrapolation may be an appropriate model for the lower part of the intertidal profile.

The quantity of data used in the extrapolation (extrapolation range) was optimised to minimise the average RMSE across all profiles. Data was used between the cut level ($X = 0$) and a second chainage further landward along the partial profile. This optimisation led to the following extrapolation ranges to be determined for each beach individually as: Gwithian (-65 m $< X < 0$ m), Wittering (-10 m $< X < 0$ m), and Brighton (-2 m $< X < 0$ m).

2.3.2. Multilayer Perceptron (MLP)

Multilayered Perceptron Models are inspired by the structure of the brain and represent one of the foundational neural network architectures (Gardner and Dorling, 1998). MLP architecture comprises multiple layers of nodes (artificial neurons) connected by weighted edges. Inputs are passed through the model layers, sequentially applying functions at each stage before outputting a result (Rana et al.,

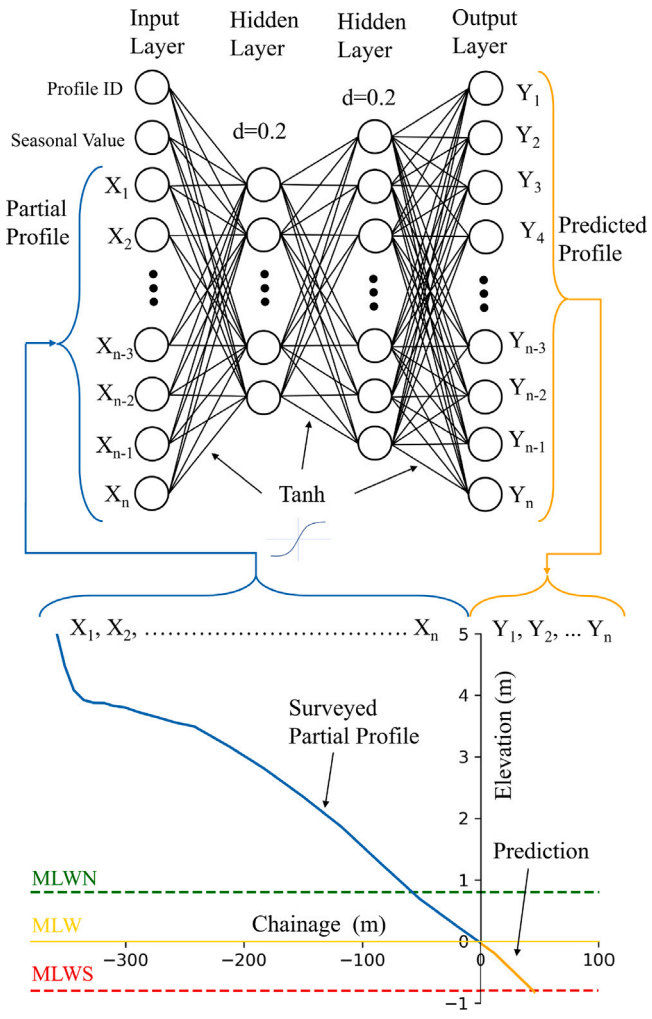

Fig. 3. MLP Network architecture, with two hidden layers. The surveyed partial profile (blue) is input into the network, alongside two profile characteristics, the profile's unique ID and a seasonal value; the prediction (orange) is the output from the network. Key hyperparameters are labelled on the diagram.

2018). This architecture has been used frequently in coastal engineering, from predicting ocean wave parameters to the damage of coastal structures (Mandal et al., 2008).

For this study, the MLP takes regularly spaced elevation data to predict the topography of the lower intertidal zone. Fig. 3 shows the network's architecture. Discretised elevations at 1 m horizontal spacing are fed into the network's input layer (X) shown in blue, with the addition of the profile characteristics discussed above (Section 2.2). The network outputs a series of elevations, again at 1 m horizontal spacing, (Y) shown in orange. The number of input and output parameters varies depending on the study location, with profiles of different lengths shown in Fig. 2.

The architecture shown includes two profile characteristics, profile ID and seasonality, which were only used in the input parameter selection tests (Section 3.1). The network's ability to easily add and remove input parameters by only altering the input layer allowed the key components of the architecture to remain unchanged, enabling the model to be trained four times with all combinations of coordinates, profile ID, and seasonal value to see the response of altering the input parameters on the prediction accuracy.

2.3.3. Convolutional Neural Networks (CNN)

Convolutional Neural Networks (CNN) are an extremely effective class of network, for which the inputs are presented in the form of matrices, where the spatial location of the data encodes additional information, such as image or video (Barnes et al., 2023). Originally inspired by the visual cortex of animals, the network performs a series of manipulations on the input matrices before passing them through fully connected layers, similar to those used in MLPs, to produce an output.

Specialising in image classification tasks, CNNs have revolutionised computer vision by enabling the extraction of meaningful information from image data. Works by Liu et al. (2021) and Ellenson et al. (2020) have both leveraged CNNs in a similar manner to investigate the temporal evolution of nearshore morphology and to classify distinct beach states and morphological features. Although typically used for classification, many CNNs have been used for regression problems, including Barnes et al. (2023) who used CNNs to improve the forecasting of total monthly regional rainfall across Great Britain. The results from all studies showed high accuracy with significant improvements compared to the current leading statistical models.

2.3.4. 2D CNN

To convert the raw survey data to images, all profiles from each location were plotted on the same axis with MLW at the origin. These images were then cropped as shown by the black box in Fig. 4 representing all of the data used to make predictions up to the cut elevation at MLW. These images were reduced to one colour channel (greyscale) and then stored as a matrix with each value representing a pixel. This process normalises the value of each pixel between 0 and 1 for greater network performance. Due to its two-dimensional structure, the 2D CNN has several orders of magnitude more input features than any of the other networks at the same study location.

The 2D CNN then takes these matrices and performs a series of transformations, manipulating them to add nonlinearity to the network. This is done through multiple different network layers seen in Fig. 4. There is a large amount of literature concerning in-depth discussion of the inner workings of a CNN, for example, Bhatt et al. (2021) provides a detailed review of current literature. However, this is beyond the scope of this paper; as such, the discussion below focuses only on the key components of the 2D CNN used.

Convolution (yellow layers in Fig. 4) is the primary transformation applied in the network. This involves passing a kernel (filter) over the matrix, multiplying the filter by the matrix and summing the result. These kernels effectively look for smaller patterns or shapes in the larger inputs, the process outputs a lower-resolution image with key features made more prominent. These kernels are not predefined; they are dynamic and are continuously updated through the training process. A large number of kernels are used on a single convolutional layer, allowing many patterns to be extracted in a single step.

Another key transformation is pooling, similar to a reduction in nodes in each sequential layer shown in the MLP network, the pooling layer reduces the spatial dimension of the subsequent layer (feature map) without the loss of important information. Max pooling is used here; the function passes over the feature map, and for each localised region, the network selects the maximum value in this area. Not only does this improve performance for small spatial translations of features, but it also significantly reduces the computation time of the network.

Batch Normalisation is the final key network layer responsible for transformations. It takes the outputs of the previous layer and normalises them over a batch of data during training. This increases the speed of the model's convergence, reducing the training time. It also introduces additional learnable parameters for each individual channel, allowing increased individuality for each kernel, resulting in improved accuracy. These layers are not shown on the diagram as they have no spatial effect on the feature maps, however, they are denoted by the Red Arrows in Fig. 4.

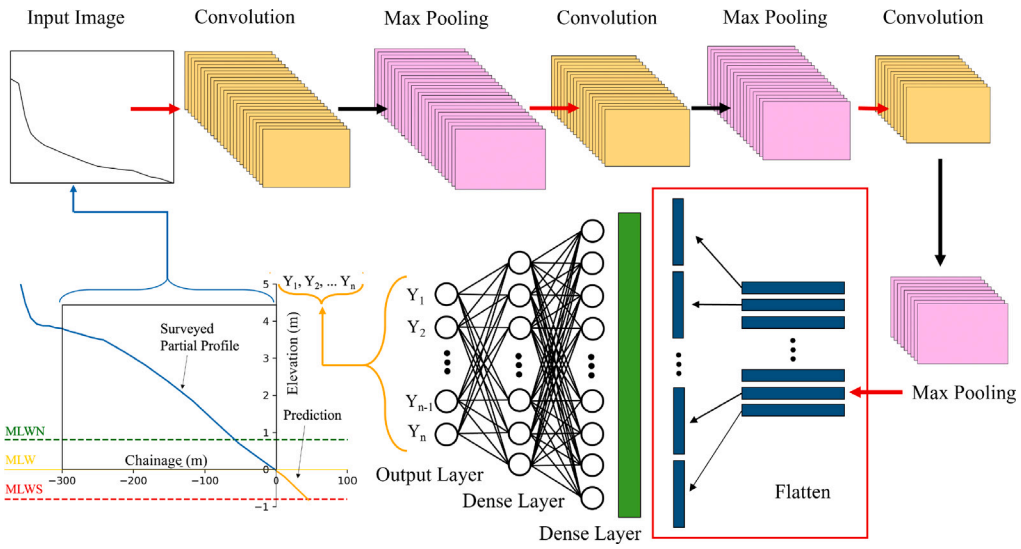

Fig. 4. 2D CNN network architecture, taking the plotted surveyed partial profile as input (black box) and outputting the predicted region (orange). The network comprises many convolution (yellow) and max pooling (pink) layers, with the red connecting arrows indicating that batch normalisation occurs at these steps. The flattening layer (red box) converts the 2D matrix into a 1D vector, which can be fed into the two fully connected dense layers that transform the predictions to the correct output size. Key hyperparameters are labelled on the diagram.

The final layers comprise fully connected dense layers as found in the MLP, which transform the output of the convolutional layers to the same cross-shore extent of the predicted region. To go from a 2D result into a 1D vector, a flattening transformation needs to occur, as seen in the red box of Fig. 4.

2.3.5. 1D CNN

The 1D CNN is a specialised case of the described 2D CNN network. While it follows the same fundamental processes, it operates on a one-dimensional input (a row vector) instead of a two-dimensional matrix. As such, instead of passing an image into the CNN, the 1 m-spaced discretised elevations (X) were passed into the network as a row vector. The length of the row vector and the output layer depended on the study location. Inputs were normalised to improve network convergence. The transformation layers remain largely the same; however, kernels become one-dimensional, with max pooling affecting only the width of the vector. As such, a flattening layer is not required, but dense layers are still used to resize the output from the network.

2.3.6. Long Short Term Memory (LSTM)

Long Short Term Memory (LSTM) networks (Hochreiter and Schmidhuber, 1997) are a type of recurrent neural network (RNN) which specialises in time series or sequential predictions. An LSTM replaces the nodes and neurons used in the MLP with memory cells and feedback loops that pass information back through the network (Joshi, 2023). With this architecture, it becomes possible to recall memories from thousands of time steps ago while selectively choosing to retain or discard information at each time step. While predominantly used in text prediction and recognition, LSTMs have been used to make time-series predictions, with many applications focusing on stock market predictions (Chen et al., 2015; Moghar and Hamiche, 2020).

As with the MLP and 1D CNN, the LSTM takes discretised partial profiles (X) shown in blue and outputs a series of predicted elevations at 1 m spacings (Y) shown in orange; the full network architecture is shown in Fig. 5. The network comprises multiple many-to-many LSTM layers connected to a dense layer that reshapes the output. The model was batch-trained and the training data shuffled due to the high number of training sequences.

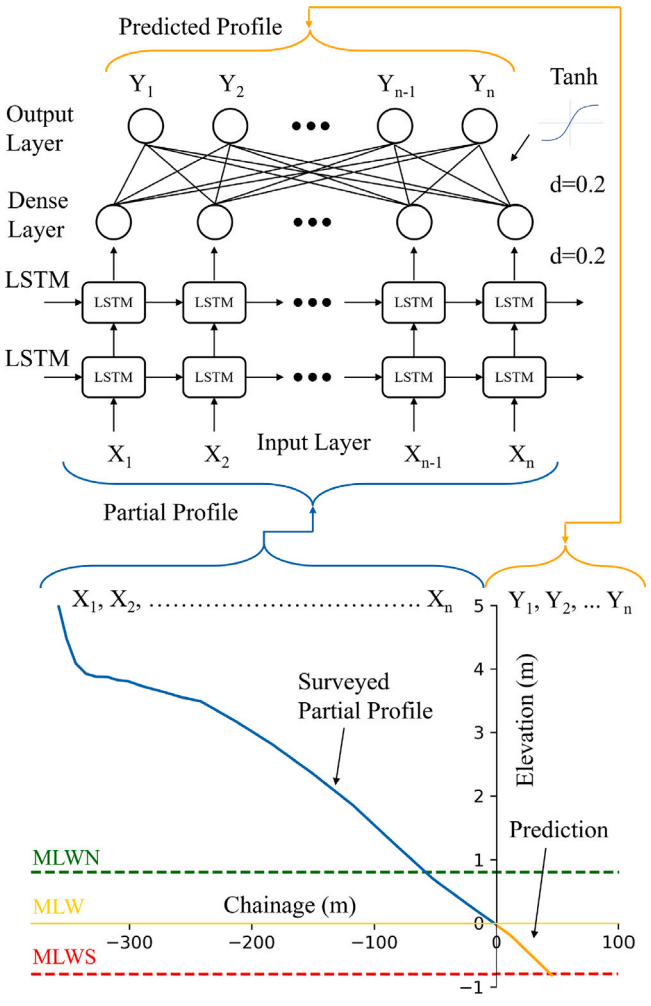

Fig. 5. Full network architecture of the LSTM, using two many-to-many model LSTM layers with an additional dense layer. Key hyperparameters are labelled on the diagram.

2.4. Hyperparameter training and optimisation

The error metric chosen to assess the network's performance was Root Mean Squared Error (RMSE). This was calculated as the difference between the discretised survey data and the predicted profile for all points below the cut elevation (MLW). Due to the limited data available, leave-one-out cross-validation (LOOCV) (a form of k-fold cross-validation) was completed to produce a more reliable performance estimate (Berrar, 2019; Lumumba et al., 2024; Weese et al., 2025). For each iteration, one profile was excluded as the test set while the model was trained and validated on the remaining profiles. This process was repeated until every profile had been used once as the test set. Statistics of RMSE were computed across all iterations to provide an unbiased estimate of predictive accuracy while maintaining the maximum available training data. A training–validation split was used to prevent overfitting, with the model trained on approximately 80% of the remaining profiles and internally validated on the remaining 20% of that subset. For the comparison of different network architectures and locations a distribution of the RMSE for each model was found in place of a single, potentially biased value of its performance, producing more effective comparisons between models.

Even with many years of training data, the number of individual profiles is still comparatively small when compared to most complex machine learning tasks. As such, maximising the available training data has a dramatic impact on the results. Batch normalisation and layer dropout were used to optimise convergence and enhance model generalisation. A dropout rate of d = 0.2 (20%) was applied to all networks to mitigate overfitting while preserving the network's learning capacity. This value is consistent with the network size and limited training data. The model was trained in batches using simple stochastic gradient descent to support an adaptive learning rate, while monitoring the validation loss. The learning rate was decreased with no improvement after 5 epochs. The training process ended when the validation loss converged and the best weights were restored. For all network architectures, the tanh activation function was selected due to its ability to preserve smooth gradients during training.

For network optimisation, hyper-band tuning was employed. Hyperband tuning is an optimisation algorithm that balances random search with a successive halving strategy. It uses effective allocation of resources to promising architectures while discarding less promising ones. It iteratively explores and evaluates a large set of hyperparameters, ultimately selecting the best-performing configuration based on the lowest value of RMSE (Li et al., 2018). All deep learning models were implemented in Python (v3.10) using the TensorFlow (v2.13) and Keras frameworks; were trained on a MacBook Pro equipped with an Intel Core i7 (2.2 GHz, 6-core) processor and 16 GB of RAM.

3. Results

3.1. Input parameter selection

To assess the impact of different input parameter selections on the network's predictive performance, various combinations of key input parameters, including elevation data, seasonality, and profile ID, were tested using the MLP. The elevation data from the partial profile made up the majority of the input data and, as such, was included in all of the tests. The seasonality and profile ID were added and removed in varying combinations. While the same logic could be applied to all neural networks, due to the MLP's structure, it has no notion of local connectivity. Local connectivity is a property of a network architecture where each neuron is connected only to a localised region, rather than to all inputs. This enables the model to focus on local patterns while simultaneously reducing the number of parameters. As such, with the MLP, additional inputs can be added easily without significant alteration to the network's architecture; it was therefore deemed the most suitable network to explore these effects.

Table 1
RMSE of each MLP with different combinations of input parameters at the three locations, presenting the mean value of the leave-one-out cross-validation across all profiles.

Model	Average RMSE (m)		
	Gwithian	Wittering	Brighton
ALL Data	0.100	0.027	0.128
Elevation	0.095	0.026	0.125
Elevation + Profile	0.096	0.026	0.125
Elevation + Seasonality	0.098	0.026	0.123

Following hyper-band tuning for network optimisation, leave-one-out cross-validation was performed on all MLP networks to evaluate the accuracy of each input combination. Each model was trained separately for each beach, and RMSE values were calculated across all relative profiles. The results are visualised using a Kernel Density Estimator (KDE) and are presented in Fig. 6. The KDE provides a non-parametric approach to assess the underlying distribution of the results, enabling comparison of model performance. For this analysis, a Gaussian kernel was used, with Scott's Rule applied to determine the kernel bandwidth. Table 1 contains the average RMSE for each model.

Fig. 6 illustrates that, although there are small differences in model accuracy, varying the input data used during training does not significantly affect performance. To ensure that seasonality and profile ID were not being overwhelmed by the large quantity of elevation data, initial weightings were added to the overall network of different magnitudes to ensure that the initial weightings of the seasonality and profile parameters were of comparable magnitudes to the sum of the elevation data. Even with a range of initial weightings, there was a negligible difference in the results.

On further inspection of Table 1, it can be seen that by adding both seasonality and profile ID, all networks performed slightly worse than networks with only elevation data as their input. The elevation-only network performed the best at both Gwithian and Wittering, indicating that there was no apparent correlation between seasonality or profile ID and the partially predicted profile elevations at this site.

In contrast to this, while the elevation-only network performed well at Brighton, the MLP that included elevation + seasonality performed the best, indicating that there may be a relationship between seasonality and predicted elevation at this site. Brighton has not only more defined beach morphology than the other study locations but also a more consistent profile shape with a low tide terrace appearing on almost every profile in the study (Fig. 2 (2)). While the elevation of a low tide terrace remains almost constant through the year, the results suggest that there is a correlation between the seasonality and the position and elevation of the low tide terrace, leading to a slight increase in network performance.

With the exception of one combination at Brighton, the analysis indicates that incorporating the supplemental data into model training at each study area does not have a significant benefit, as evidenced by increased RMSE values for all locations. Therefore, the subsequent analysis will rely solely on elevation data as input to the networks.

3.2. Model comparison

Four different neural networks, MLP (Section 2.3.2), 2D CNN (Section 2.3.4), 1D CNN (Section 2.3.5), and LSTM (Section 2.3.6), were compared to a baseline Linear extrapolation (Section 2.3.1) model to determine the optimum network architecture for this application.

To assess the ability of different networks to predict partial beach profiles, hyper-band tuning for network optimisation, followed by leave-one-out cross-validation, was again performed on all networks. As with the input parameter combinations, the models for each study location were trained separately. The results of the cross-validation are

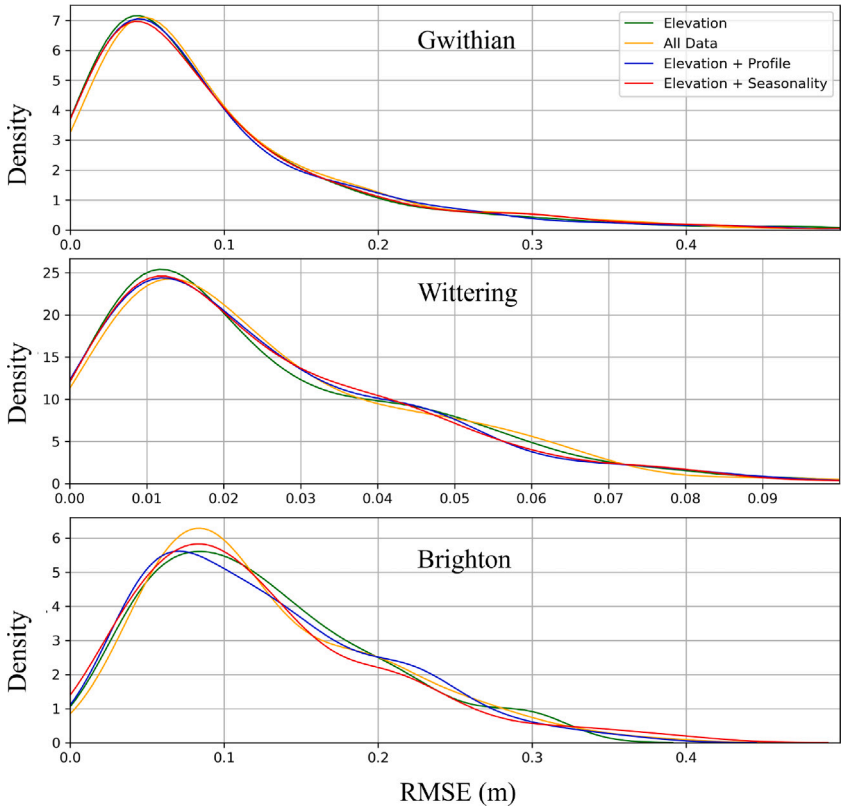

Fig. 6. Kernel density estimation of RMSE for the four MLPs with varying input parameters. Shown for all three study locations Gwithian, Wittering, and Brighton.

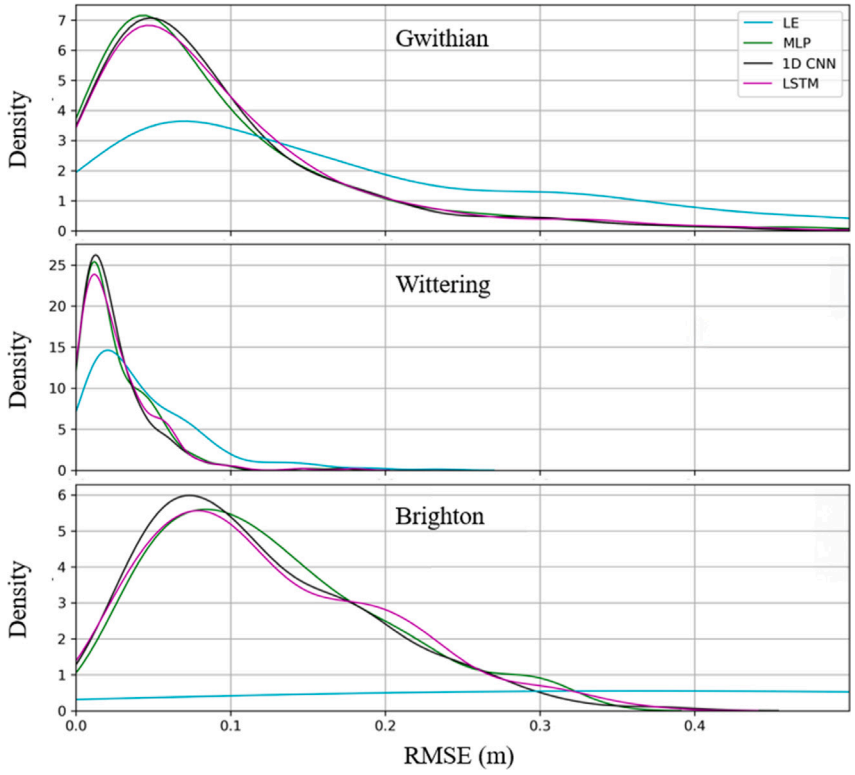

Fig. 7. Kernel Density Estimation for the leave-one-out cross-validation results is shown for all five predictive models at each location, Gwithian, Wittering, and Brighton.

Table 2
Mean RMSE performance of each model across the three locations, based on leave-one-out cross-validation. The 2D CNN failed to converge during training and therefore did not produce meaningful results.

Model	Average RMSE (m)		
	Gwithian	Wittering	Brighton
LE	0.180	0.044	0.989
MLP	0.095	0.026	0.125
2D CNN	NAN	NAN	NAN
1D CNN	0.096	0.025	0.119
LSTM	0.098	0.026	0.124

presented using a Kernel Density Estimator (KDE) (Fig. 7). The mean values of the RMSE for each model are presented in Table 2.

All models were able to predict partial beach profiles, except the 2D CNN, which was unable to converge during training, regardless of network architecture or parameter tuning. Even with a much more complex structure than the other networks, there was not enough complexity or training data to prevent under-fitting. As such, regardless of the input elevations, the network predicted the same profile (within a few millimetres) which minimised the total error. The network's complexity also resulted in a much higher computational time. As such, it was excluded from all subsequent analyses.

The results show that the success of each predictive model varied substantially depending on the beach type. The remaining four models were most effective for Wittering, with an average RMSE of between 0.025–0.044 m. The models were not only more accurate but also more consistent with much narrower frequency kernels. The models had interquartile ranges (IQRs) of: MLP 0.029 m, LSTM 0.027 m, 1D CNN 0.023 m, and LE 0.043 m. In contrast, Brighton produced the most inconsistent results with IQRs between 0.093 and 0.964 m; all models struggled to consistently predict beach profiles at this location.

As expected, the baseline model (Linear Extrapolation) was associated with the greatest error, with significantly higher RMSE than all neural networks for all three locations. This is especially evident at Brighton, with an average RMSE almost 8 times greater than the worst-performing neural network. From Fig. 2 (2) it can be seen that the beach profile at Brighton is highly non-linear with a low tide terrace (LTT) on almost all profiles starting just above MLWS. As seen in Fig. 8 (e,f) a LTT exhibits a sudden large change in gradient seaward of the cut point and so a linear fit is a poor approximation, resulting in a greatly diverging prediction and large error. In contrast to this, LE has much better accuracy at Wittering (Fig. 8(c)), where the profile is almost linear, leading to an RMSE larger, but of the same magnitude as the neural networks.

The MLP and LSTM performed largely similarly, with average RMSE values differing by between 0.00–0.003 m across all beaches. They both consistently predicted profiles to a high level of accuracy, with the second-best performance across all models. While there was only a small difference in RMSE between the networks, there was a large difference between the profiles that they predicted. The MLP profiles are very irregular with large jumps in elevation between consecutive points, as shown in Fig. 8 (a,b,e,f). The predicted profile often oscillates above and below the real profile, highlighted in Fig. 8(e). The LSTM profiles are significantly smoother with gradual changes between points, replicating the real-world characteristics of beaches. As such, the LSTM is considered significantly better at predicting realistic beach profiles. While MLPs offer many advantages, their architecture does not account for spatial or sequential relationships between neighbouring input parameters or output values. The network has no local connectivity, and as such, changing the order in which the input data is passed into the network has no impact, as long as this is kept consistent through the training process. An LSTM is a type of recurrent neural network, and as such, the order of input and output parameters is important, as it has a high notion of local connectivity. The network assumes that for all points, the points before and after are important, with every adjacent elevation value being related. The network architecture adheres to series predictions, as seen by the pink profiles in Fig. 8, the network produces much smoother profiles and more realistic results.

Comparing the performance of the MLP and LSTM at Brighton, while the LSTM can predict the presence and shape of the LTT, due to its sequential nature, the network's lack of spatial awareness along the profile often results in the predicted LTT being at the incorrect elevation. The MLP performed similarly in this regard. The MLP predicted profiles that predominantly included the LTT in the correct position; however, due to the issues previously discussed, the predicted shape was often noisy (Fig. 8 e,f). Additionally, it is noted that in some measured profiles there is no LTT present or the LTT is at a very different elevation to the majority of the data (Fig. 8(f)). Although the MLP has spatial awareness within the profile, there is no information relating to the position of the profile in a 3D space. The MLP and LSTM were unable to predict the aberrant profiles, and this resulted in the second and third points of inflection in the KDE in Fig. 7 (MLP: inflections at RMSE = 0.19 m and 0.28 m; LSTM: inflections at RMSE = 0.18 m and 0.30 m).

Many of the profiles at Gwithian featured bar morphology, as observed in Fig. 8(b), resulting in the largest range of elevations seaward of MLW. Both networks struggled to accurately predict the occurrence or location of bar morphology, often resulting in profiles with high RMSE. Many of the Gwithian profiles exhibited only small bar formations with minimal elevation change. The LSTM was able to reproduce reasonably realistic-looking bar geometry, as seen in pink in Fig. 8(b), but the magnitude and location of the predicted bars typically did not match reality and sometimes led to large divergence from the measured profile. While the MLP did not predict any significant bar morphology, the resultant profile often bisected the bar (Fig. 8(b)), resulting in a lower RMSE than the LSTM.

The best performing network was the 1D CNN, with the lowest RMSE across two of the three study locations and the second lowest at Gwithian. Unlike the other networks that had clear secondary points of inflection at higher RMSE's, at all three locations the CNN had a single peak at a lower RMSE and a steep decline at higher errors (Fig. 7).

The profiles predicted by the 1D CNN are not as smooth as those produced by the LSTM, but are significantly less noisy than the MLP profiles. The 1D CNN uses kernels to extract information from the input data by assessing patterns between adjacent elevations. This introduces local connectivity within the network. In contrast to this, the last few layers of the network are dense layers to transform the output of the convolutional layers into the correct size to make predictions. As discussed for the MLP, the dense layers are fully connected and, as such, have no local connectivity. The LSTM has a single dense layer, while the 1D CNN has multiple, decreasing the local connectivity in the network.

3.3. Morphological features

The presence of morphological features was found to be the dominant factor affecting accuracy when predicting beach profiles. Profiles with high morphological nonlinearity (i.e. large perturbations from a linear profile) exhibited significantly higher errors than those with fewer features.

Quantifying the scale of features on a topographic profile is challenging, with many varying definitions. The two properties used in this study are profile curvature and morphological nonlinearity. Profile curvature is the mean of the second derivative across the predicted portion of the profile ($f(x)$). The second metric, morphological nonlinearity, is the standard deviation of elevation change in the linearly detrended profile. This is used to quantify large deviations between

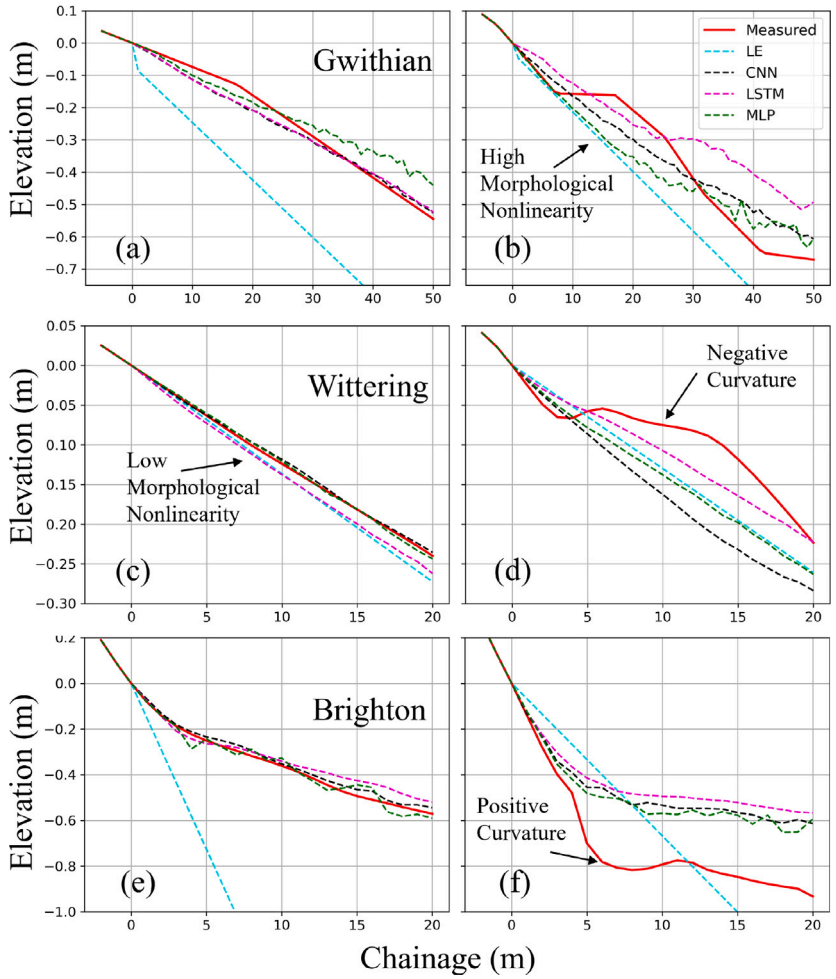

Fig. 8. Examples of predicted profiles produced by all four models, shown alongside the corresponding measured profile in Red. The left column (a,c,e) shows profiles predicted well by the neural networks. The right column (b,d,f) shows instances where the neural networks were unable to predict the profile with any accuracy. (a) Profile at Gwithian displaying a gentle curve with minimal morphological features. (b) Profile displaying bar morphology in the predicted region of the profile, poorly predicted by all models. (c) Typical featureless profile at Wittering. (d) Poorly predicted profile at Wittering due to anomalous bar morphology on this profile. All plots demonstrate that profiles with aberrant morphological features result in predictions with the highest error. (e) Profile with low tide terrace at c. −0.4 m elevation, predicted well as the LTT often occurs at this elevation. (f) Low tide terrace, predicted poorly due to the significantly lower elevation than most of the profiles at Brighton.

adjacent points present in rapidly changing geometry. It is detrended to ensure that profile steepness has no impact on this measure.

$$\text{Curvature} = \frac{1}{L}\int_0^L f''(x)\,dx \tag{2}$$

$$\text{Morphological Nonlinearity} = \frac{\sigma(\Delta E)}{m} \tag{3}$$

Where L is the profile length, $\sigma(\Delta E)$ is the standard deviation of the elevation differences between adjacent points (ΔE), and m is the linear gradient of the profile. Examples of both curvature and morphological nonlinearity are labelled in Fig. 8.

Seen in Fig. 9(a), profile curvature is plotted against prediction error for the MLP, 1D CNN and LSTM models for all three study locations. Profiles that have a negative curvature are associated with a local maximum, such as sand bar morphology. Positive curvatures are caused by morphological features which create local minima, such as an LTT (Brighton), long-shore bar troughs (Gwithian), or beach steps (Brighton).

For both Gwithian and Wittering, the RMSE is smallest for profiles with zero curvature (see Fig. 9). Negative curvature is observed at both Gwithian and Wittering, as a result of large bar systems at Gwithian and smaller bars at Wittering and causes a substantial increase in RMSE. Similarly, positive curvature at these locations results in a similar increase in RMSE. Some profiles at Gwithian show evidence of a bar trough in the predicted region, and Wittering has an occasional LTT, resulting in positive curvature. The most prominent morphological features that produce high curvature magnitudes at both Gwithian and Wittering are not frequently represented in the training data and are therefore prone to overfitting on relatively featureless profiles.

At Brighton, the RMSE is largely independent of curvature, with only a small increase observed as the magnitude of curvature increases (see Fig. 9). A low tide terrace is present in nearly all Brighton profiles, as such the curvature in this case does not describe the presence of varying morphological features but the position or slope of the LTT. At Brighton, increased positive curvature indicates profiles where the LTT slopes gradually upwards (Fig. 8(e) 4 m < X < 12 m) rather than flattening completely. The network's ability to make accurate predictions largely depends on the cross-shore position and elevation of the LTT it is trying to predict, and is independent of the LTT slope; therefore, RMSE remains constant with curvature.

Fig. 9(b) clearly shows increasing RMSE with morphological nonlinearity at all study locations. Sudden morphological changes, as indicated by high values of morphological nonlinearity, are difficult to predict and are uncommon in the profile data. Consequently, the models either overfit relatively featureless profiles or fail to capture the

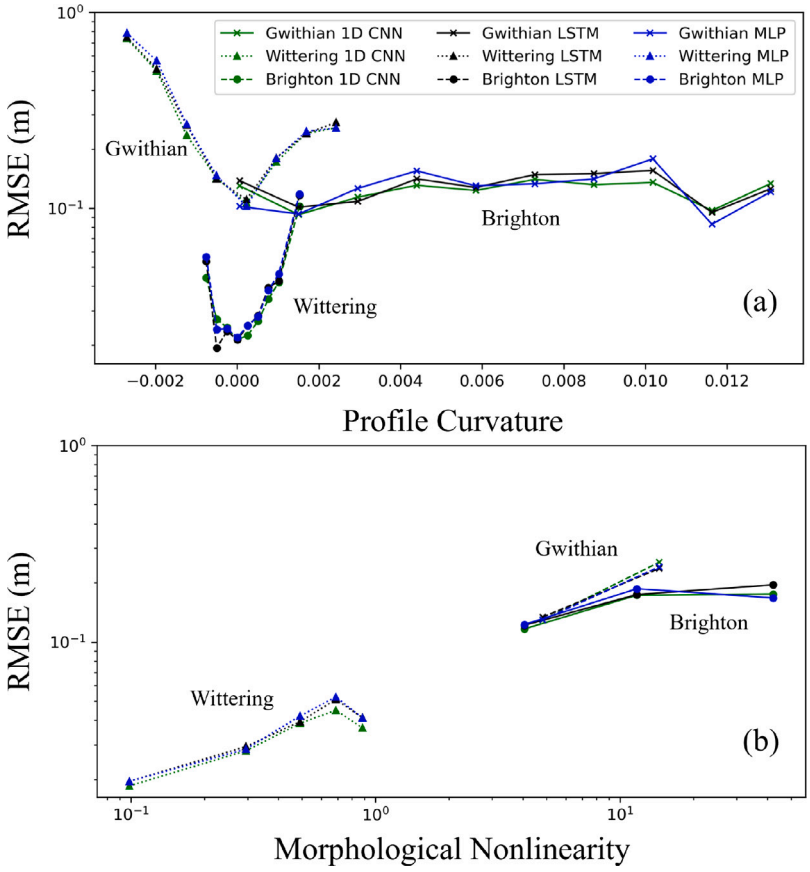

Fig. 9. Average RMSE value at all three study locations for the 1D CNN model, binned by (a) Profile Curvature and (b) Profile Nonlinearity.

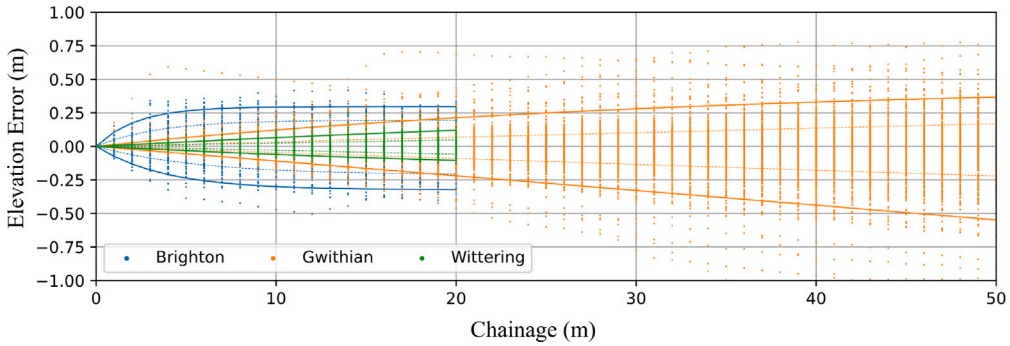

Fig. 10. Spatial distribution of prediction error by the 1D CNN model for all three study locations. Confidence intervals of 15%/75% (dashed line) and 5%/95% (solid line) are marked for each location.

complex relationships. As a result, all networks at all locations struggle to accurately predict profiles with large morphological nonlinearity.

As the 1D CNN was the best performing network, the following analysis was conducted on the results of the 1D CNN. To assess the spatial variability of the error produced by the network, the error from each predicted profile is plotted as a function of its chainage. While maximum error can be easily observed, bootstrap confidence intervals have been added to relate elevation error to its probability of occurrence. Fig. 10 presents the elevation error as a function of chainage for each study location, including upper and lower confidence intervals.

At Wittering, the elevation error at both the upper and lower bound confidence intervals increases linearly with chainage, reflecting Wittering's linear beach gradient. At Brighton, the cross-shore variability of elevation error is primarily a result of incorrect predictions of the LTT elevation. The elevation error initially increases cross-shore (0 m < x < 7 m) as the partial predicted profile and the surveyed profiles diverge, up to the landward limit of the LTT (x = 7 m). As both the predicted and measured low tide terraces have almost the same gradient, the error remains relatively constant throughout the rest of the profile. As such, the confidence intervals take an exponential form. The elevation error at Gwithian increases with chainage in a near-linear fashion, exhibiting a slight quadratic curvature. This subtle non-linearity is attributed to the bar morphology characteristic of Gwithian.

Network interpretability is pivotal for building trust in these neural networks, transforming them from black-box models into reliable tools. While the inner workings of neural networks are complex, explainability allows us to understand how and why a network arrived at its prediction. Without this, it is hard to validate results and any potential bias of predictive models. These techniques can also provide valuable insights and highlight hidden trends in the underlying data.

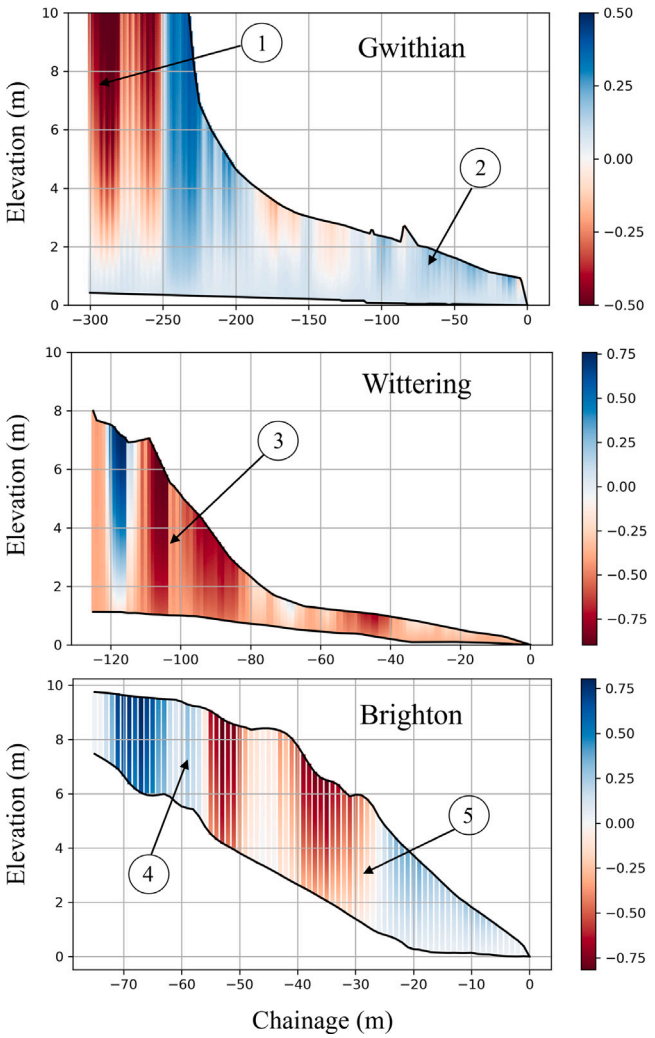

Fig. 11. Results of a Monte Carlo simulation performed on the 1D CNN. The colour represents the impact of each coordinate on the average elevation of the predicted profile. Only coordinates within the surveyed beach extents (black lines) are used in the simulation.

A Monte Carlo simulation was performed on the 1D CNN to assess the impact of each coordinate (Chainage, Elevation) within the measured beach envelope on the predicted partial profile. This was done by predicting the profile using only a single coordinate within the beach envelope of the measured partial profiles (black line, Fig. 11), building on and adapting the methods of Barnes et al. (2023) and Sundararajan et al. (2017). Each coordinate was input into the network individually, with all other values set to 0. The predicted partial profile was then averaged to produce a single value representing its average elevation. This value was compared to the average elevation of the mean surveyed profile in the same region to determine whether that coordinate contributed to an increase or decrease in profile elevation. The results are presented as a colourmap in Fig. 11. Using this method, we can begin to understand the relationship between the surveyed partial profile input to the network and the model's predicted output.

Using Fig. 11, it can be seen that if a surveyed partial profile passes through a red region of the graph, the resulting predicted partial profile will have a lower average elevation. Conversely, if the profile passes through a blue region, this will increase the average elevation of the predicted partial profile. White areas indicate regions that either have no impact on the prediction and are deemed unimportant by the neural network or correspond to profiles with an average elevation similar to the mean.

The way in which the colour varies with elevation at a single cross-shore position can be used to understand the importance of this location on the prediction. For example, areas of the map that have the same colour, regardless of elevation, indicates a consistent influence on the prediction, regardless of the elevation at those points. As such, this region is unimportant to the overall predicted profile. In contrast, areas with large vertical colour gradients are locations where the elevation has the most impact on the predicted partial profile. Some interesting regions of Fig. 11 are discussed below and labelled within the figure.

(1) The largest influence at Gwithian originates from the landward end (between x = −300 m and x = −250 m) of the measured profile, with this region contributing up to a 0.5 m decrease in the average elevation of the predicted profile per point. The significant colour gradient indicates that this location is important for making predictions, with higher elevations here causing a dramatic decrease in the average elevation of the predicted region. By contrast, at the lower parts of the beach (2) (between x = −100 m and x = 0 m), higher elevations result in an increase in the elevation of the predicted partial profile.

The results for Wittering are largely consistent, with all locations seaward of −100 m causing a nearly identical reduction in predicted profile elevation regardless of the input elevation in this region. As such, the predicted profiles are almost entirely predicted using profile data between x = −120 m and x = −100 m marked (3) on Fig. 11. This is partially attributed to the largest variation in surveyed profiles in this region (widely spaced black lines). This region contains two clear vertical colour gradients, one blue (between x = −122 m and x = −114 m) and one red (between x = −115 m and x = −104 m), both showing decreasing values at lower elevations. If the surveyed partial profiles passes through this region at a steeper angle intersect both a darker blue and a lighter red zone, resulting in a flatter or higher average elevation in the partial predicted profile. Flatter surveyed partial profiles through this region will pass through both a dark blue and dark red, cancelling each other out, resulting in a lower average elevation of the partial predicted profile.

Brighton displays the opposite behaviour to Wittering and Gwithian, with the landward part of the beach ($x < -45$) having little impact on the predicted partial profile. Region (4) exhibits no vertical colour gradient, indicating that the elevation of the surveyed profile in this area does not influence the predicted profile. Instead, all elevations at these chainages contribute equally to the prediction. Unlike the other two study locations, the surveyed beach envelope does not narrow greatly before reaching the MLW level (0 m). As such, there is enough variation in the beach toe to have a significant impact on the predicted region. Shown in region (5) (between x = −45 m and x = 0 m), the greater the elevation of the surveyed profile in this location, the greater the average elevation of the predicted profile.

4. Discussion

While the error of each model has been quantified, it is important to understand the implications of these errors for coastal engineering and coastal management. In the UK, the main uses of beach profiles by coastal engineers and managers, outside of scientific study, are beach volume calculations and wave overtopping or coastal flood risk assessments (Harley et al., 2011).

To measure the effects of using partially predicted profiles for volumetric calculations, Fig. 12(a-c) presents a rain plot of the percentage difference between the volumes calculated using predicted and measured profiles. For these calculations, the volume of each profile has been calculated for a single metre width of the beach. Upper and lower elevation constraints were implemented to calculate beach volume within a defined vertical range.

Beach profile volumes are used primarily to monitor the net sediment transport over time to assess the health of a beach and any trend of erosion or accretion. Over shorter time scales, beach volumes can experience cyclic seasonal variations that are a response to a

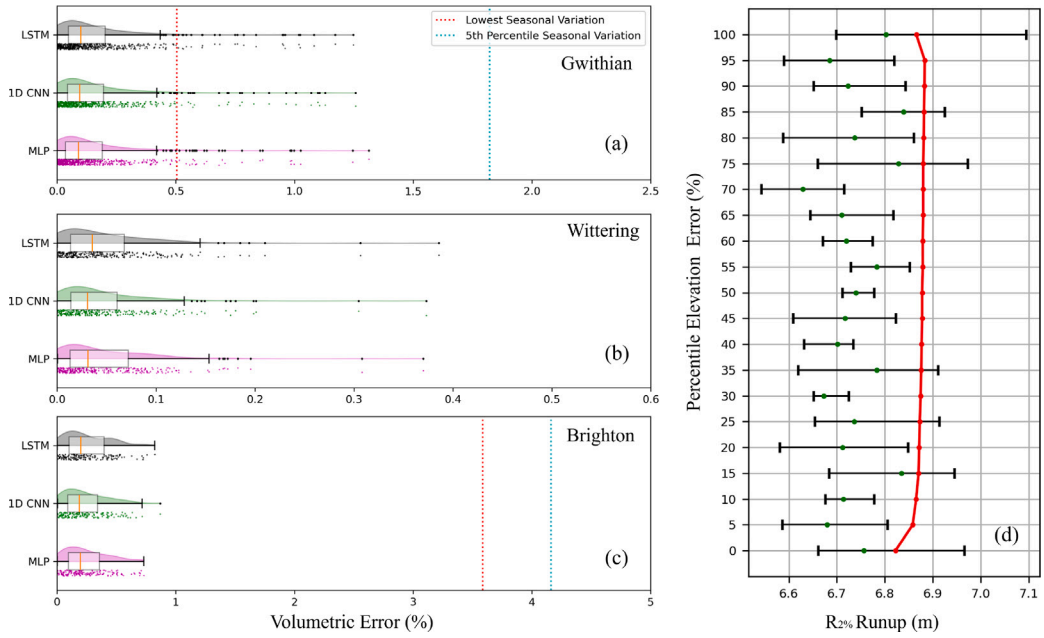

Fig. 12. (a-c) A rain plot showing the percentage error in beach volume calculations when using predicted profiles, plotted for each model at each study location. Also shown is the seasonal variation in beach volume at each location, with the red line indicating the lowest recorded variation during the study period and the blue line marking the 5th percentile of seasonal variation. (d) Variation in runup resulting from applying a range of elevation errors to the same beach profile. These errors are defined by the confidence intervals calculated from the 1D CNN. $R_{2\%}$ is shown for both an identical sequence of irregular waves (red line) and for four simulations per profile using irregular waves (JONSWAP spectrum) with a different seeding value (black error bars, with the average indicated in green).

varying wave climate throughout the year. To calculate the seasonal variation at each study location, the volume of all profiles surveyed more than once during a single year was calculated, and the variation between profiles was assessed. Where possible, both the 5th percentile of seasonal variation and the lowest recorded value are plotted on the graph, shown in blue and red respectively. Wittering (Fig. 12(c)) had a minimum seasonal variation of 3.89% and a 5th percentile of 13.28%.

The seasonal variation of beach volume at all study sites is significantly higher than the volumetric error produced by using partially predicted profiles. As such, it can be assumed that the error introduced will have a minimal effect on the assessment of seasonal variation in beach profile volumes and indeed any effect is comparable to the error introduced by other factors related to survey practice. For example, Silveira et al. (2013) demonstrated that the longshore interval between profiles had a larger impact on volumetric calculations for their study site, where doubling the profile spacing from 100 to 200 m led to a 2%–3% change in the calculated beach volume. This is larger than the error introduced by using partially predicted profiles for beach volume calculations.

To measure the impacts of using partially predicted profiles for wave overtopping prediction, XBeach was used to simulate wave run-up ($R_{2\%}$) for profiles with different magnitudes of elevation error. While Xbeach has demonstrated its ability to model the physics of overtopping (McCall et al., 2010), there are limitations on its ability to model coastal flooding on natural beaches (Stokes et al., 2021). The ability of XBeach to model wave runup is more robust, and as such, the elevation exceeded by 2% of wave runup events $R_{2\%}$ was selected as the metric used to assess the effect of errors in the predicted partial profile (de Beer et al., 2021).

To assess changes in runup due to elevation errors, a baseline initial profile was used as a reference, and varying magnitudes of error were introduced into the predicted region. The baseline chosen was the mean profile at Gwithian (Blue line, Fig. 2 Gwithian). To introduce error that reflected the error introduced by the neural network, the confidence intervals from the 1D CNN were chosen (displayed in orange in Fig. 10). While only four are shown in Fig. 10, these confidence intervals were calculated for 5th percentile increments from 5%–95%, with the 0th and 100th percentiles representing both the maximum and minimum error in predictions made by the network. These confidence intervals were then applied to the predicted region of the baseline profile to generate a set of new profiles with varying magnitudes of error, reflective of the network's predictions.

A storm condition of H_s = 3 m, T_p = 8 s was simulated using a surfbeat model, at a water level of mean high water springs (MHWS) for a 1-hour period. Two sets of simulations were completed for each modified profile: (1) an identical sequence of irregular waves by choosing an identical seeding value, and (2) four repeats of irregular waves (JONSWAP spectrum) for each profile with a different seeding value. Fig. 12(d) indicates that there are no significant changes in the predicted values of $R_{2\%}$ for varying profile error seaward of MLW (x = 0 m). In the case of identical irregular wave series, the value of $R_{2\%}$ is almost independent of profile elevation error seaward of x = 0 m, varying by only 0.0215 m for percentage errors between 0 and 100%. For simulations with the same wave spectrum but varying seeding values, the black bars represent the spread in runup across the four repeated tests for each profile. It can be seen that the variations in run-up predictions caused the wave sequence generated by a fixed spectrum are significantly higher than the change in runup predictions caused by the alterations in morphology.

As most simulations are completed with a spectral wave input, it can be said that the error in predicted R2% introduced by errors in the partially predicted profile would not have any significant effect on wave run-up or overtopping modelling.

5. Conclusion

In this study, five predictive models (LE, MLP, 2D CNN, 1D CNN, and LSTM) were used to predict the lower intertidal zone between MLW and MLWS using partial beach profiles on three beaches of different types. Linear extrapolation, used as the baseline model, performed the worst, with an average RMSE between 0.044–0.989 m. The LSTM and MLP models were significantly better with an average RMSE between 0.026–0.125 m, while the 1D CNN model performed the best with an average RMSE of between 0.025–0.119 m across the three study

locations. It was found that the network's notion of local connectivity, the idea of each neuron only being connected to a localised region, played a pivotal role in producing realistic beach profiles. Networks with high local connectivity, like the LSTM, produced smooth profiles, while the MLP, with no local connectivity, produced profiles with rapidly varying elevations not present in the smooth surveyed profiles.

The observation that different neural network architectures achieve similar predictive accuracy suggests that the dataset itself sets an upper limit on performance. Despite having very different inductive biases, MLPs treat all inputs as equally related without considering spatial or temporal structure, CNNs exploit local spatial patterns and correlations among nearby points, and LSTMs capture sequential dependencies, no model was able to surpass this level of accuracy. This implies that the information contained in the dataset is inherently limited, and further improvements will likely require additional or higher-quality data rather than more complex architectures.

The key factor in the accuracy of predicted profiles was the cross-shore morphological nonlinearity at each location, with beaches containing fewer features having a lower RMSE. The models performed best at Wittering due to its featureless nature, but struggled most at Brighton, which had features (LTT) on almost every profile. This trend was further seen within the study locations, with individual profiles having a higher curvature and morphological nonlinearity displaying higher values of RMSE.

The 1D CNN model could be a useful tool for beach managers who use topographic beach data to make decisions about our coastlines. The elevation error introduced by using neural networks as a predictive tool has been shown to have a negligible effect when using this data for volumetric and overtopping analysis. All models were trained on a MacBook Pro equipped with an Intel Core i7 (2.2 GHz, 6-core) processor and 16 GB of RAM. Full hyperparameter tuning was completed in under 24 h per location, demonstrating that the approach is computationally efficient and feasible on mid-range hardware.

The method is, however, subject to certain practical constraints. To enable the network to learn the complex patterns necessary for predicting partial profiles, a large volume of training data is required, necessitating the prior collection of a substantial dataset. As such, this approach may not be suitable for locations without an established coastal monitoring programme. The method may not be suitable for highly dynamic coastlines undergoing significant long-term morphological change. Substantial changes in beach morphology at a location may necessitate reinstating full survey extents in order to retrain the model and maintain predictive accuracy. Furthermore, as previously noted, while the method is well-suited for coastal monitoring applications, it does not offer the level of accuracy required for coastal process studies.

This method has demonstrated its value in mitigating information loss if the seaward extent of topographic surveys was changed from MLWS to MLW. As such, it may be possible to revise the existing extents of the survey specification, with baseline surveys (conducted at significantly reduced frequency) still extending down to the current level of MLWS. This would support ongoing model validation and provide additional training data to help mitigate errors associated with long-term morphological change. This approach would not only reduce the pressure on the current survey schedule but also allow more frequent surveys/additional locations with the same resources.

CRediT authorship contribution statement

Samuel Rose: Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Chris Blenkinsopp:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Conceptualization. **Andrew Barnes:** Conceptualization, Methodology, Supervision, Validation, Writing – review & editing. **William Russell:** Writing – review & editing, Data curation. **Charlie Thompson:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Barnes, A.P., McCullen, N., Kjeldsen, T.R., 2023. Forecasting seasonal to sub-seasonal rainfall in Great Britain using convolutional-neural networks. Theor. Appl. Climatol. 151 (1), 421–432. http://dx.doi.org/10.1007/s00704-022-04242-x.

Berrar, D., 2019. Cross-validation. In: Ranganathan, S., Gribskov, M., Nakai, K., Schönbach, C. (Eds.), Encyclopedia of Bioinformatics and Computational Biology. Academic Press, Oxford, pp. 542–545. http://dx.doi.org/10.1016/B978-0-12-809633-8.20349-X.

Bhatt, D., Patel, C., Talsania, H., Patel, J., Vaghela, R., Pandya, S., Modi, K., Ghayvat, H., 2021. CNN variants for computer vision: History, architecture, application, challenges and future scope. Electronics 10 (20), 2470. http://dx.doi.org/10.3390/electronics10202470.

Bruun, P., 1954. Coast Erosion and the Development of Beach Profiles. Report, United States, Beach Erosion Board, URL https://erdc-library.erdc.dren.mil/jspui/handle/11681/3426.

Castelle, B., Masselink, G., 2023. Morphodynamics of wave-dominated beaches. Camb. Prism.: Coast. Futur. 1, e1. http://dx.doi.org/10.1017/cft.2022.2.

Channel Coastal Observatory, 2021. National coastal monitoring. URL https://coastalmonitoring.org/.

Chen, K., Zhou, Y., Dai, F., 2015. A LSTM-based method for stock returns prediction: A case study of China stock market. In: 2015 IEEE International Conference on Big Data (Big Data). pp. 2823–2824. http://dx.doi.org/10.1109/BigData.2015.7364089.

de Beer, A.F., McCall, R.T., Long, J.W., Tissier, M.F.S., Reniers, A.J.H.M., 2021. Simulating wave runup on an intermediate–reflective beach using a wave-resolving and a wave-averaged version of XBeach. Coast. Eng. 163, 103788. http://dx.doi.org/10.1016/j.coastaleng.2020.103788.

Dean, R.G., 1977. Equilibrium beach profiles: U.S. atlantic and gulf coasts. In: Ocean engineering technical report, Dept. of Civil Engineering and College of Marine Studies, University of Delaware, Newark, Del..

Dean, R.G., 1991. Equilibrium beach profiles: Characteristics and applications. J. Coast. Res. 7 (1), 53–84, URL https://www.jstor.org/stable/4297805.

Ellenson, A.N., Simmons, J.A., Wilson, G.W., Hesser, T.J., Splinter, K.D., 2020. Beach state recognition using argus imagery and convolutional neural networks. Remote. Sens. 12 (23), 3953. http://dx.doi.org/10.3390/rs12233953.

Fenneman, N.M., 1902. Development of the profile of equilibrium of the subaqueous shore terrace. J. Geol. 10 (1), 1–32, Publisher: The University of Chicago Press. URL https://www.jstor.org/stable/30055541.

Gardner, M.W., Dorling, S.R., 1998. Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. Atmos. Environ. 32 (14), 2627–2636. http://dx.doi.org/10.1016/S1352-2310(97)00447-0.

Harley, M.D., Turner, I.L., Short, A.D., Ranasinghe, R., 2011. Assessment and integration of conventional, RTK-GPS and image-derived beach survey methods for daily to decadal coastal monitoring. Coast. Eng. 58 (2), 194–205. http://dx.doi.org/10.1016/j.coastaleng.2010.09.006.

Hashemi, M.R., Ghadampour, Z., Neill, S.P., 2010. Using an artificial neural network to model seasonal changes in beach profiles. Ocean Eng. 37 (14), 1345–1356. http://dx.doi.org/10.1016/j.oceaneng.2010.07.004.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Comput. 9 (8), 1735–1780. http://dx.doi.org/10.1162/neco.1997.9.8.1735.

Joshi, A.V., 2023. Machine Learning and Artificial Intelligence. Springer International Publishing, Cham, http://dx.doi.org/10.1007/978-3-031-12282-8.

Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., Talwalkar, A., 2018. Hyperband: A novel bandit-based approach to hyperparameter optimization. http://dx.doi.org/10.48550/arXiv.1603.06560, arXiv:1603.06560 [cs].

Li, Y., Van Oosterom, P., Ge, Y., Zhang, X., Baart, F., 2020. A CNN–LSTM method for the morphology evolution prediction of beach mega-nourishment. IEEE Access 8, 184512–184523. http://dx.doi.org/10.1109/ACCESS.2020.3030119.

Liu, B., Yang, B., Masoud-Ansari, S., Wang, H., Gahegan, M., 2021. Coastal image classification and pattern recognition: Tairua beach, New Zealand. Sensors 21 (21), 7352. http://dx.doi.org/10.3390/s21217352.

López, I., Aragonés, L., Villacampa, Y., Compañ, P., 2018a. Artificial neural network modeling of cross-shore profile on sand beaches: The coast of the province of Valencia (Spain). Mar. Georesources Geotechnol. 36 (6), 698–708. http://dx.doi.org/10.1080/1064119X.2017.1385666.

López, I., Aragonés, L., Villacampa, Y., Satorre, R., 2018b. Modelling the cross-shore beach profiles of sandy beaches with Posidonia oceanica using artificial neural networks: Murcia (Spain) as study case. Appl. Ocean Res. 74, 205–216. http://dx.doi.org/10.1016/j.apor.2018.03.004.

Lumumba, V.W., Kiprotich, D., Mpaine, M.L., Makena, N.G., Kavita, M.D., 2024. Comparative analysis of cross-validation techniques: LOOCV, K-folds cross-validation, and repeated K-folds cross-validation in machine learning models. Am. J. Theor. Appl. Stat. 13 (5), 127–137. http://dx.doi.org/10.11648/j.ajtas.20241305.13, Publisher: Science Publishing Group.

Mandal, S., Patil, s., Manjunatha, Y., Hegde, A., 2008. Application of neural networks in coastal engineering - An overview. In: 12th International Conference on Computer Methods and Advances in Geomechanics 2008, vol. 3.

Masselink, G., Short, A.D., 1993. The effect of tide range on beach morphodynamics and morphology: A conceptual beach model. J. Coast. Res. 9 (3), 785–800, Publisher: Coastal Education & Research Foundation, Inc. URL https://www.jstor.org/stable/4298129.

McCall, R.T., Van Thiel de Vries, J.S.M., Plant, N.G., Van Dongeren, A.R., Roelvink, J.A., Thompson, D.M., Reniers, A.J.H.M., 2010. Two-dimensional time dependent hurricane overwash and erosion modeling at Santa Rosa Island. Coast. Eng. 57 (7), 668–683. http://dx.doi.org/10.1016/j.coastaleng.2010.02.006.

Moghar, A., Hamiche, M., 2020. Stock market prediction using LSTM recurrent neural network. The 11th International Conference on Ambient Systems, Networks and Technologies (ANT) / The 3rd International Conference on Emerging Data and Industry 4.0 (EDI40) / Affiliated Workshops, Procedia Comput. Sci. The 11th International Conference on Ambient Systems, Networks and Technologies (ANT) / The 3rd International Conference on Emerging Data and Industry 4.0 (EDI40) / Affiliated Workshops, 170, 1168–1173.http://dx.doi.org/10.1016/j.procs.2020.03.049,

NNRCMP, 2023. National coastal monitoring - specifications and scopes. https://coastalmonitoring.org/ccoresources/specificationsandbriefs/. (Accessed 8 April 2023).

Phillips, B.T., Brown, J.M., Bidlot, J.R., Plater, A.J., 2017. Role of beach morphology in wave overtopping hazard assessment. J. Mar. Sci. Eng. 5 (1), 1. http://dx.doi.org/10.3390/jmse5010001.

Phillips, B.T., Brown, J.M., Plater, A.J., 2020. Modeling impact of intertidal foreshore evolution on gravel barrier erosion and wave runup with XBeach-X. J. Mar. Sci. Eng. 8 (11), 914. http://dx.doi.org/10.3390/jmse8110914.

Rana, A., Singh Rawat, A., Bijalwan, A., Bahuguna, H., 2018. Application of multi layer (perceptron) artificial neural network in the diagnosis system: A systematic review. In: 2018 International Conference on Research in Intelligent and Computing in Engineering. RICE, pp. 1–6. http://dx.doi.org/10.1109/RICE.2018.8509069.

Saville, Jr., T., Caldwell, J.M., 1952. Accuracy of hydrographic surveying in and near the surf zone. Coast. Eng. Proc. (3), http://dx.doi.org/10.9753/icce.v3.3, 3–3.

Scott, T., Masselink, G., Russell, P., 2011. Morphodynamic characteristics and classification of beaches in England and Wales. Mar. Geol. 286 (1), 1–20. http://dx.doi.org/10.1016/j.margeo.2011.04.004.

Silveira, T., Carapuço, M., Sousa, H., Taborda, R., Psuty, N., Andrade, C., Freitas, M., 2013. Optimizing beach topographical field surveys: Matching the effort with the objectives. J. Coast. Res. 65, 588–593. http://dx.doi.org/10.2112/SI65-100.1.

Stokes, K., Poate, T., Masselink, G., King, E., Saulter, A., Ely, N., 2021. Forecasting coastal overtopping at engineered and naturally defended coastlines. Coast. Eng. 164, 103827. http://dx.doi.org/10.1016/j.coastaleng.2020.103827.

Sundararajan, M., Taly, A., Yan, Q., 2017. Axiomatic attribution for deep networks. In: Proceedings of the 34th International Conference on Machine Learning. PMLR, (ISSN: 2640-3498) pp. 3319–3328.

Vellinga, P., 1982. Beach and dune erosion during storm surges. Coast. Eng. 6 (4), 361–387. http://dx.doi.org/10.1016/0378-3839(82)90007-2.

Weese, M.L., Smucker, B.J., Edwards, D.J., 2025. The use of cross validation in the analysis of designed experiments. http://dx.doi.org/10.48550/arXiv.2506.14593, arXiv:2506.14593 [stat].

Wright, L.D., Short, A.D., 1984. Morphodynamic variability of surf zones and beaches: A synthesis. Mar. Geol. 56 (1), 93–118. http://dx.doi.org/10.1016/0025-3227(84)90008-2.