Offshore Technology Conference 2024

OTC-35104

A spatiotemporal machine learning framework for the prediction of metocean conditions in the Gulf of Mexico

Edward Steele[1], Jiaxin Chen[2,3], Ian Ashton[3], Ajit Pillai[3], Sergio Jaramillo[4], Pak Leung[5] and Luz Zarate[5]

[1]Met Office, FitzRoy Road, Exeter, EX1 3PB, UK
[2]University of Plymouth, Drake Circus, Plymouth, PL4 8AA, UK
[3]University of Exeter, Penryn Campus, Penryn, TR10 9FE, UK
[4]Shell Global Solutions US Inc., 150 North Dairy Ashford Road, Houston, TX 77079, USA
[5]Shell International Exploration and Production Inc., 3333 Highway 6 South, Houston, TX 77082, USA

Abstract

Machine learning techniques offer the potential to revolutionize the provision of metocean forecasts critical to the safe and successful operation of offshore infrastructure, leveraging the asset-level accuracy of point-based observations in conjunction with the benefits of the extended coverage (both temporally and spatially) of numerical modelling and satellite remote sensing data. Here, we adapt and apply a promising framework – originally proposed by the present authors for the prediction of wave conditions on the European North West Shelf – to the waters of the Gulf of Mexico. The approach consists of using an attention-based long short-term memory recurrent neural network to learn the temporal patterns from a network of available buoy observations, that is then combined with a random forest based spatial nowcasting model, trained on reanalysis data, to develop a complete framework for spatiotemporal prediction for the basin. By way of demonstration, the new method is applied for the short-range prediction of wave conditions up to 12 hours ahead, using in-situ wave observations from the sparse network of National Data Buoy Center locations as an input, with the corresponding spatial mapping learned from the physics-based Met Office WAVEWATCH III global wave hindcast. The full spatiotemporal forecast system is assessed using independent measurements in the vicinity of the Louisiana Offshore Oil Port, previously unseen by the machine learning model. Results show that accurate real-time, rapidly updating wave predictions are possible, available at a fraction of the computational cost of traditional physics-based methods. The success of the approach, combined with the flexibility of the framework, further suggest its utility in related metocean challenges. While still at an early stage of development into a fully relocatable capability, it is intended that this contribution provides a foundation to stimulate a series of subsequent efforts to help support improved offshore planning and workability – including (but not limited to) applications linked with better resolving spatial variability across renewable energy sites, predicting ocean current regimes in the proximity of oil & gas platforms, as well as informing adaptive sampling strategies conducted by autonomous vessels – where the adoption of such a machine learning approach, that can be run on a laptop computer, having the potential to revolutionize data-driven decision-making by the industry.

Introduction

We are entering an exciting new era of data-driven weather prediction, where forecast models trained on historical data (including observations and reanalyses) offer an alternative to directly solving the governing equations of fluid dynamics (Schultz et al., 2021). By capitalizing on the vast amount of available historical data – and capturing their inherent patterns that are not easily represented via explicit equations – such machine learning weather prediction (MLWP) techniques have the potential to increase both forecast accuracy and efficiency in comparison to traditional numerical weather prediction (NWP) equivalents, where such information is more difficult to distil and the cost of computation is more expensive (Lam et al., 2023). Indeed, recent years have seen a succession of celebrated advances in the development of global MLWP models of the atmosphere – including Keisler (2022), Bi et al. (2022), Pathak et al. (2022) and Lam et al. (2023) – that have reported comparable performance to leading operational NWP models for a selection of variables at 1.0° and 0.25° resolution; highlighting the emergent opportunities afforded by these types of approaches.

Motivated primarily by the need to support the safe and successful operation of offshore infrastructure, the complementary development of regional MLWP models of the ocean have also been the subject of increasing research interest (Sonnewald et al.

2021; Song et al., 2023). While the use of machine learning techniques – loosely defined as the collection of algorithms (e.g. support vector machines, SVMs, random forests, RFs, and artificial neural networks, ANNs) for solving multi-variate, non-linear, non-parametric classification or regression problems – are themselves not new, recent advances in deep learning architectures have catalyzed their wider application, including forecasting of fundamental metocean parameters such as surface currents (e.g. Muhamed-Ali et al., 2021; Sinha & Abernathey, 2021) and waves (e.g. James et al., 2018; Fan et al. 2020). Popular deep learning architectures include multi-layer perceptrons (MLPs; comprising multiple sequential layers of neurons, each connected to the neurons of the previous and next layers), convolutional neural networks (CNNs; comprising multiple filter layers that weight its inputs, producing so-called 'feature maps'), and recurrent neural networks (RNNs; comprising a hidden state to store information about the history of the sequences presented to its inputs; particularly useful in time-series forecasting). Reviews of some of the ocean-based applications of these approaches are presented in Sonnewald et al. (2021) and Song et al. (2023), however despite the similarity in the algorithms used, the potential for providing a standard, scalable and relocatable framework – that is both region and parameter agnostic – has not yet been realized.

In this paper, we adapt and apply the promising Machine Learning for Low-Cost Offshore Modelling (MaLCOM) framework – originally proposed by the present authors for the prediction of wave conditions on part of the European North West Shelf (Chen et al., 2021, 2023) – to the waters of the Gulf of Mexico (GoM). The approach consists of using an attention-based long short-term memory (LSTM) RNN to learn the temporal patterns from a network of available observations, combined with an RF-based spatial nowcast model trained on reanalysis data, to develop a complete method for spatiotemporal forecasting for the entire region. Within a smaller and simpler area, Chen et al. (2023) reported equivalent performance to a leading operational NWP model under typical conditions, based on three locations as an input. In the present study, we are therefore motivated to test the approach within a different area, more than 30 times larger in size, under conditions of extreme sparsity, to better inform future developments. Here, forecasts are made using in-situ wave observations from six National Data Buoy Center (NDBC) buoy locations as an input, with the corresponding spatial mapping learned from the physics-based Met Office WAVEWATCH III global wave hindcast at 25 km resolution, and the performance assessed using independent measurements in the vicinity of the Louisiana Offshore Oil Port, previously unseen by the machine learning model. The structure and flexibility of the framework suggests its potential further utility in related metocean challenges, such as in enabling enhanced prediction of complex Loop Current (LC) and Loop Current Eddy (LCE) dynamics. While still at an early stage of refinement, it is intended that this contribution provides the basis for stimulating a series of subsequent efforts to directly support improved offshore planning and workability – including (but not limited to) applications linked with better resolving spatial variability across renewable energy sites, predicting ocean current regimes in the proximity of oil & gas platforms, as well as informing adaptive sampling strategies conducted by autonomous vessels – where the adoption of such a machine learning approach, that can be run on a laptop computer, has the potential to revolutionize data-driven decision-making by the industry.

The structure of this paper is as follows: the 'Models and Methods' section presents an overview of the machine learning approach that forms the basis of the framework developed, while the 'Application' section presents an example of its use for short-range forecast purposes. The final section presents a discussion of the concept, as well as outlining opportunities for further work.

Models and Methods

The complete framework consists of combining a site-specific temporal forecasting model (Chen et al., 2023) with a spatial nowcasting model (Chen et al., 2021), as shown in Figure 1. Here, the key steps include the pre-processing of the real-time input data, the machine learning-based mapping from these observations to the spatial forecast, and the post-processing/visualization of the output data. In full spatiotemporal prediction mode, the machine learning component uses an attention-based LSTM-RNN to produce a temporal forecast of the conditions at each of the observation locations, that in turn force an RF-based spatial model driven by the conditions at these same locations. The temporal model was trained using two years of historical wave buoy observations for each site, while the spatial model was trained using two years of physics-based hindcast wave data, with subsequent testing on a further unseen month, as described below.

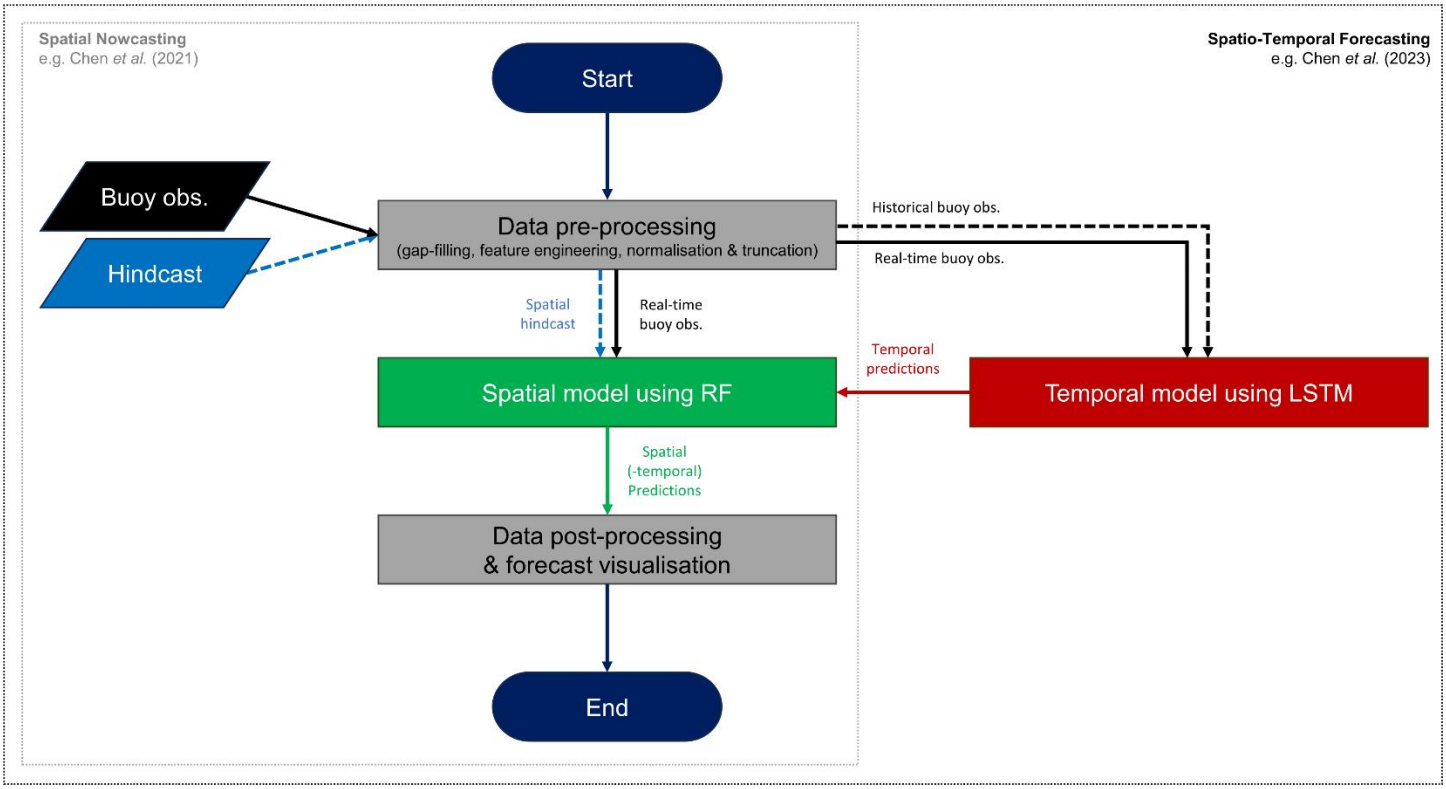

Figure 1: Overview of the MaLCOM framework, where the key steps include the pre-processing of the real-time input data, the machine learning-based mapping from these observations to the spatial forecast, and the post-processing/visualization of the output data. In full spatiotemporal prediction mode, the machine learning component uses an attention-based LSTM-RNN to produce a temporal forecast of the conditions at each of the observation locations, that in turn force an RF-based spatial model driven by the conditions at these same locations. Here, the dashed lines indicate the data used for model training, while the solid lines indicate the data used in operational real-time mode.

Figure 2 presents a map of the bathymetry of the GoM, corresponding to the area of the model domain. The colored markers indicate the position of the six input wave buoy locations (red dots), as well as the independent evaluation site (blue star), with the data collected by a combination of Datawell Directional Wave Rider Mk III and NDBC Self Contained Ocean Observing Payload (SCOOP) platforms. Coincident observations of significant wave height (H_s), mean wave direction (M_{Dir}) mean wave period (T_m), peak wave period (T_p) were recorded for all six input wave buoy locations during the period 1 January 2021 to 31 January 2023, with training conducted for the period 1 January 2021 to 31 December 2022, and subsequent testing (evaluation) conducted for the unseen month of January 2023.

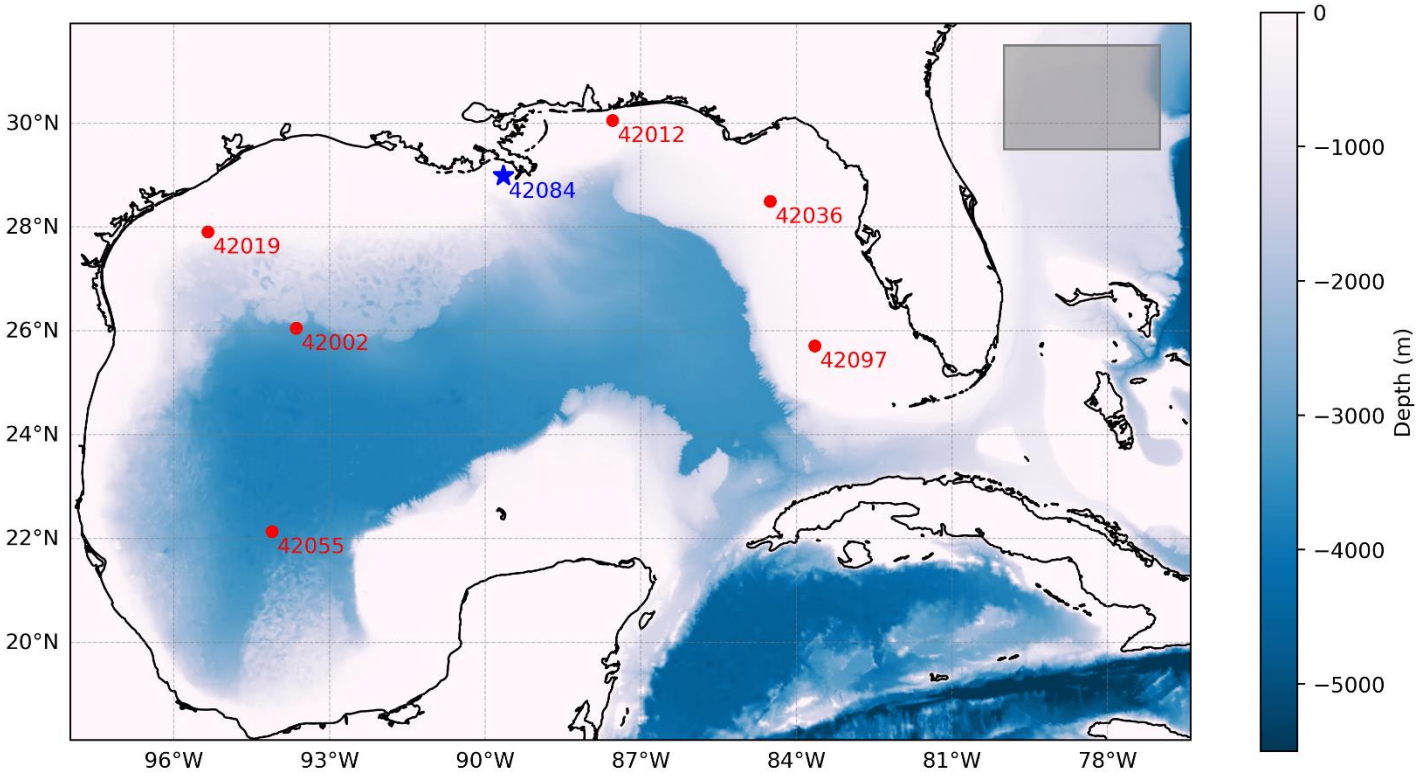

Figure 2: Map of the bathymetry of the GoM, corresponding to the area of the model domain. The colored markers indicate the position of the six input wave buoy locations (red dots), as well as independent evaluation site (blue star). Here, the gray box in the top right hand corner denotes the size of the area originally used in the MaLCOM framework development of Chen et al. (2021) and Chen et al. (2023).

The pre-processing of the observations is particularly important for the training of the temporal model. As with all methods involving timeseries prediction, gap filling is necessary to maintain the continuity and coherence of the observations, as well as maximizing the quantity of data able to be used and minimizing the potential for skewing the statistics. Since the present in-situ records are largely complete (i.e. mostly only requiring alignment onto a consistent half-hourly time base) then this is achieved via linear interpolation of each of the four parameters from each of the six input wave buoy locations independently. However, as metocean parameters are known to exhibit seasonal, annual and interannual trends, it is not necessarily always appropriate to simply rely on indiscriminate temporal patterns to learn the future behavior. To suitably represent this seasonality, multiple training datasets are therefore constructed using a multi-year rolling three-month window (within which waves are typically well correlated) for subsequent application based on the forecast period sharing the same month on which it is centered. To prevent high-magnitude variables (such as wave direction) from overwhelming low-magnitude variables (such as significant wave height), which may then make the model unstable and suffer from poor performance and increased sensitivity to large input values, all observations are rescaled to have an interval of [0, 1], using a maximum-minimum normalization process (thus requiring the output of the temporal model to later be transformed back to absolute values using linear rescaling). To further support the optimal selection of forcing variables, an elastic net algorithm (Pirhooshyaran & Snyder, 2020) with cross-validation (Meyer et al., 2018) was employed to choose the best features for temporal model training, resulting in the inclusion of all of the parameters that were available from the observations (i.e. H_s, M_{Dir}, T_m and T_p). Finally, the training data was truncated into shorter sequences for the temporal model, using a sliding window, consistent with its multistep lookback / multistep forecast structure. Note that unlike with deep learning models, there is no requirement for the inputs to RF models to undergo any engineering, normalization or transformation of their features, owing to the differences in how these data are inherently handled (Molnar, 2020; Chen et al., 2021).

The temporal model is fully described by Chen et al. (2023). This is based on an attention encoder-decoder LSTM-RNN structure, a type of sequence-to-sequence timeseries model, used to forecast conditions at multiple future time steps according to conditions at multiple previous time steps. For the present study, forecasts are generated at a half-hourly interval up to 12 hours (i.e. 24 time steps) ahead, considering the previous 24 hours i.e. (48 time steps) of observations, consistent with findings by Pirhooshyaran et al. (2020) that a lookback period of twice the length of the forecast horizon is optimal with respect to accuracy. Compared with a simple encoder-decoder LSTM-RNN structure (Sutskever et al., 2014) configured to read the input sequence, encode it, decode it and transfer it to further layers – but importantly only providing the hidden state from the last time step from the encoder – the attention mechanism (Luong et al., 2015) exposes the output from both the encoder and the decoder at each timestep, helping the

network identify the parts of the input that are more correlated with the target elements of a prediction task and therefore improving the retention ('memory') of information from earlier timesteps (Du et al., 2020). Following Chen et al. (2023), we additionally incorporate a time-distributed dense layer between the encoder–decoder layer and the output, further increasing the complexity and potentially the accuracy of the network; with the evaluation of the performance of the whole model based on its ability to regress the output sequence.

The spatial model is fully described by Chen et al. (2021). This is based on an RF structure, a type of ensemble learning regression model, in which multiple parallel decision trees are used to infer values following a hierarchy of binary rules determined from training data (Breiman, 2001). For the present study, the output from the physics-based Met Office WAVEWATCH III global hindcast model was used to define the mapping between the individual wave buoy locations and the rest of the model domain as described by the coordinates of [18.12°N to 33.88°N] and [98.00°W to 77.00°W]. This hindcast is forced by atmospheric 10 m wind fields from the European Centre for Medium Range Weather Forecasts (ECMWF) ERA5 Global Reanalysis at 0.25 degree resolution (Hersbach et al., 2020) and configured on a three-tier Spherical Multiple-Cell (SMC) grid mesh (Li, 2012; Saulter et al., 2016), comprising 25 km-12 km-6 km resolution cells that are then regridded to a regular 20 km grid of 60 x 60 grid points for the convenience of application. The wave model parameterizations follow those of the other Met Office operational systems (Valiente et al., 2023), with a spectral resolution of 36 linearly-spaced directional bins and 30 logarithmically-spaced frequency bins covering a period of between 25 and 1.5 seconds, from which integrated parameters, equivalent to the observations, including H_s, M_{Dir}, T_m and T_p can be extracted.

Application

The application of the spatiotemporal machine learning approach is demonstrated for the short-range prediction of wave conditions up to 12 hours ahead in the GoM.

Figure 3 presents an example spatial comparison of the significant wave height (H_s) output of the physics-based hindcast ('Hc') model with that of the machine learning ('ML') model, relative to the 12Z run on 1 January 2023. Here, the wave height from the hindcast model is shown in the left hand column, the wave height obtained from the machine learning model is shown in the middle column, and the difference between the two (calculated as ML-Hc) is shown in the right hand column, at a lead time of T+0 (top row), T+6 (middle row) and T+12 (bottom row) hours ahead. Consistent with Figure 2, the colored markers indicate the position of the six input wave buoy locations (red dots), as well as independent evaluation site (blue star). Encouragingly, despite the extreme sparsity of these input data, it is seen that the machine learning model impressively resolves the spatial structure of the wave fields, typically within ±25 cm of the hindcast model, particularly when mapping between adjacent input wave buoy locations with similar site characteristics (e.g. depth and exposure). Where larger differences in wave height of up to ±75 cm occur, then this is typically in areas that are less well constrained/represented by the input observations and their expected spatial translation, recognizing the distances between these as well as large bathymetric range in the basin. While the temporal evolution of the forecast appears stable, it is noted there is a small increase in the resulting spatially-averaged mean absolute error (MAE) from 0.1677 m to 0.1840 m, and root mean squared error (RMSE) from 0.0444 m to 0.0477 m at lead times of T+0 and T+12 hours ahead, respectively. A further assessment of the machine learning model performance is considered in Figure 4.

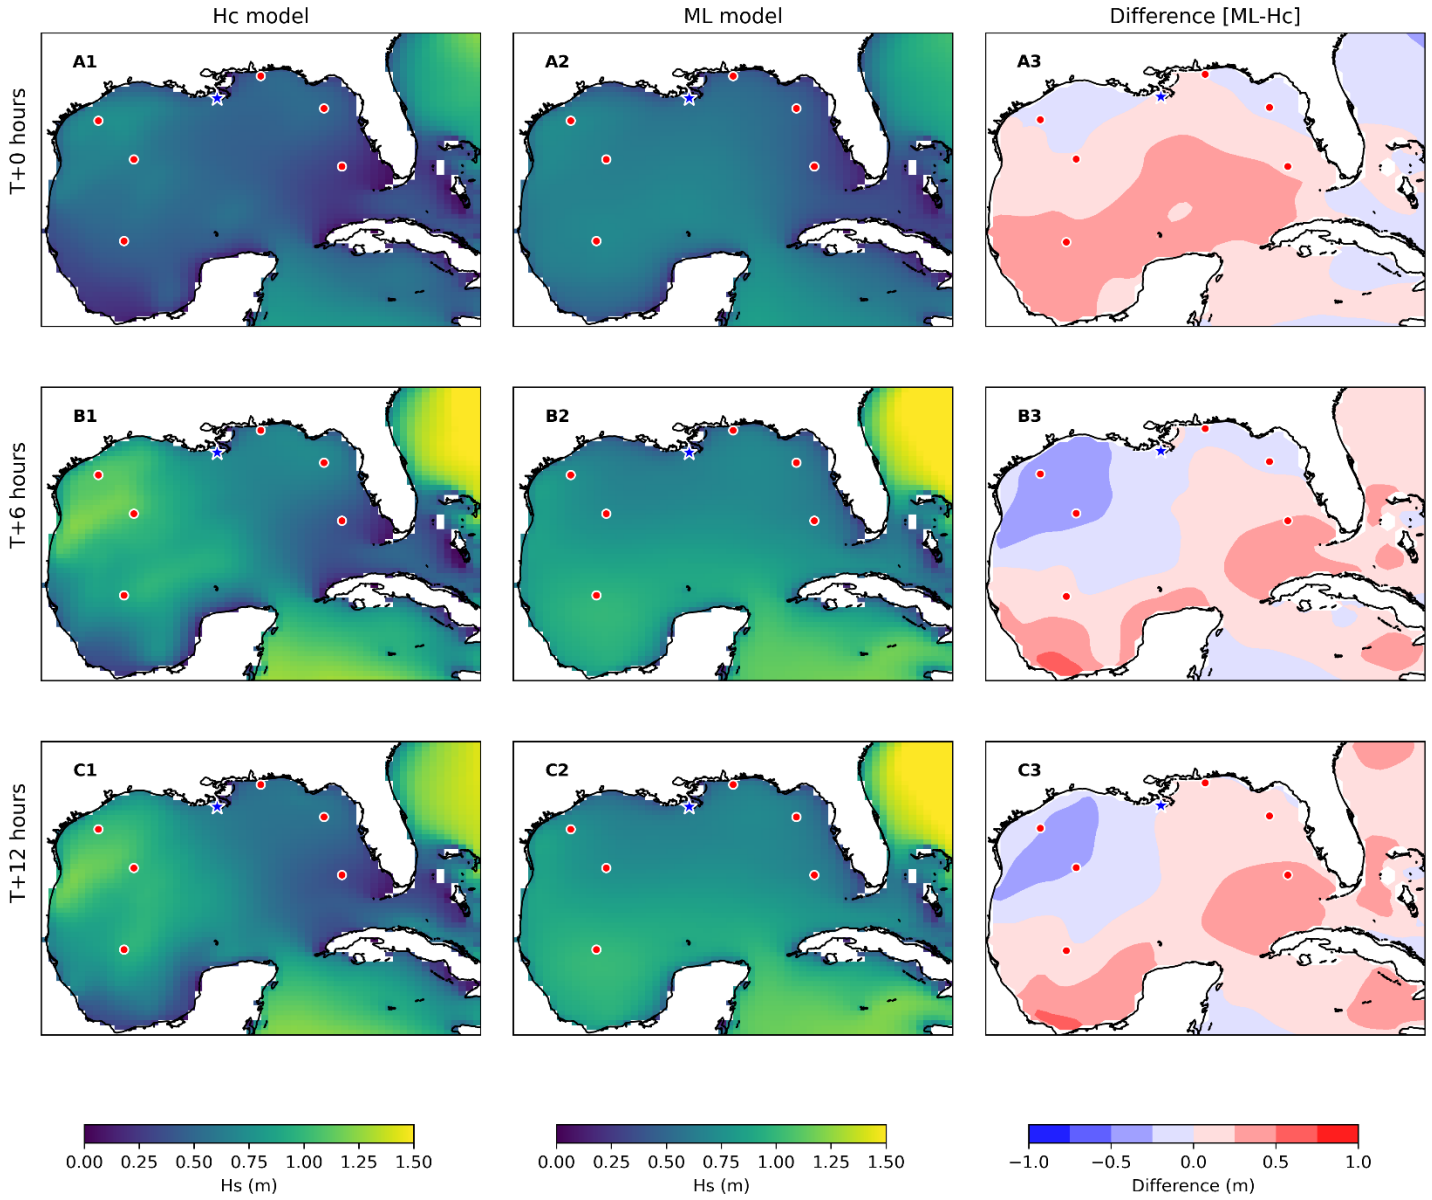

Figure 3: Example spatial comparison of the significant wave height (H_s) output of the physics-based hindcast ('Hc') model with that of the machine learning ('ML') model, relative to the 12Z run on 1 January 2023. The wave height from the hindcast model is shown in the left hand column, the wave height obtained from the machine learning model is shown in the middle column, and the difference between the two (i.e. ML-Hc) is shown in the right hand column, at a forecast lead time of T+0 (top row), T+6 (middle row) and T+12 (bottom row) hours ahead. Consistent with Figure 2, the overlaid markers indicate the six input wave buoy locations (red dots) and independent evaluation site (blue star).

Figure 4 presents an example temporal comparison of the significant wave height (H_s) output of the machine learning ('ML') model with that of the independent observations from the NDBC wave buoy ("buoy obs.") in the vicinity of the Louisiana Offshore Oil Port, using one month of data from January 2023. This site (marked by the blue star in the earlier figures and previously unseen by the model) is near the head of the Mississippi Canyon, 55 km southeast of Port Fourchon, Louisiana. Significantly, it is 235 km from the nearest of the six input wave buoy locations, and on the far side of the Mississippi Riva Delta which (together with its coastal characteristics) combine to make it more challenging to predict. The timeseries of the T+0 (i.e. nowcast) performance (black line) is shown in Figure 4A, relative to that of the hindcast (solid blue line) and the buoy observations (dashed blue line). Here, it is seen the machine learning model suitably resolves the temporal structure of the in-situ wave data, albeit with a tendency to over-estimate the wave height consistent with that of the hindcast data on which the spatial model was trained (e.g. 1-5 January and 21-25 January, for example). This difference is particularly prominent during periods of low wave height (e.g. $H_s < 0.5$ m) and likely associated with the extreme sparsity of the inputs adversely affecting its ability to reconstruct the spatial field (see Appendix 1), as well as the quantity of previous events with similar combinations of wave characteristics seen within the training period, resulting in a lower coefficient of determination (r^2) of 0.64 compared to 0.86 (see Appendix 2). While limited in opportunities to test this

further due to the constraints of available in-situ data, previous work (Chen et al., 2021, 2023) suggests that increasing both the density of input observations and the length of the training period would increase skill, achieving parity of performance compared to traditional physics-based methods. To complement this assessment, scatterplots of the machine learning model forecast performance at a lead time of T+0 (left hand column), T+6 (middle column) and T+12 (right hand column) hours ahead are shown in Figure 4B, where the first of these is simply an alternative presentation of the same comparison presented in Figure 4A. As expected, it is noted there is a small decrease in performance with forecast horizon, largely arising from the accuracy of the prediction of the future H_s, M_{Dir}, T_m and T_p conditions at the input wave buoy locations, indicated by a decrease in the coefficient of determination (r^2) from 0.64 to 0.55 at lead times T+0 and T+12 hours ahead, respectively. However, again, this is expected to be able to be remedied by increasing the length of the training period.

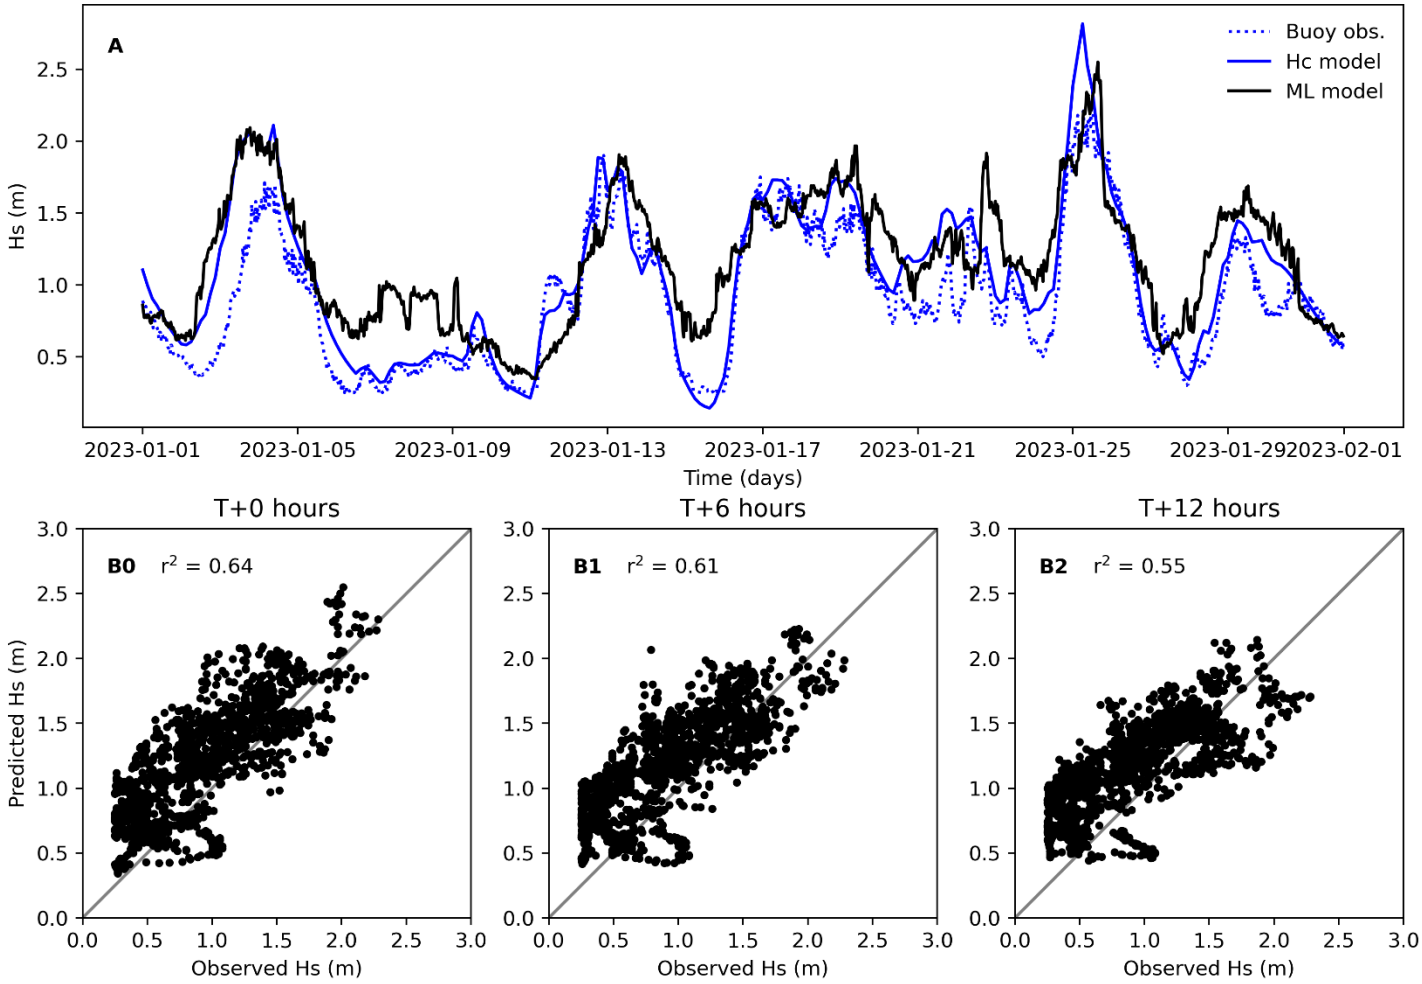

Figure 4: Example temporal comparison of the significant wave height (H_s) output of the machine learning ('ML') model with that of the independent observations from the NDBC wave buoy ("buoy obs.") in the vicinity of the Louisiana Offshore Oil Port, using one month of data from January 2023. Figure 4A: Timeseries of the machine learning model performance ('ML', black line), relative to that of the physics-based hindcast ('Hc', solid blue line) and the buoy observations ('Buoy obs.', dashed blue line) at a nowcast lead time of T+0. Figure 4B: Scatterplots of machine learning model performance, relative to the buoy observations, at a lead time of T+0 (left hand column), T+6 (middle column) and T+12 (right hand column) hours ahead. The diagonal gray line shows an (ideal) 1:1 relationship.

Discussion

These results show the real potential for the application of a spatiotemporal machine learning approach, combining an LSTM-RNN and an RF-based surrogate model, for the improved prediction of metocean conditions in the GoM. Once trained, the persistence model can be run at low computational cost, taking advantage of both rapidly updating observations (which enable the model to issue more frequent forecasts), and existing regional physics-based hindcast data (which enable the model to achieve higher spatial resolution) (Chen et al., 2023). While opportunities for the further development of the GoM configuration are acknowledged, this is deemed an important tool in supporting improved offshore planning and workability – including (but not limited to) applications linked with better resolving

spatial variability across renewable energy sites, predicting ocean current regimes in the proximity of oil & gas platforms, as well as informing adaptive sampling strategies conducted by autonomous vessels – where the adoption of such a machine learning approach, that can be run on a laptop computer, has the potential to revolutionize data-driven decision-making by the industry.

As with all machine learning techniques, the accuracy of the output is contingent on the availability of suitable data from which to learn the relevant spatial and temporal patterns. For the present study, only two years of data (from 1 January 2021 to 31 December 2022) were used to train both the LSTM-RNN and the RF model components, based on the conditions at only six wave buoy locations (to which we were constrained by the very limited available/continuous wave observations reporting H_s, M_{Dir}, T_m and T_p). Compared with the previous application of the framework for the prediction of wave conditions within a 49,000 km^2 area of the European North West Shelf – that reported equivalent performance to leading operational NWP models under typical conditions – the 1,600,000 km^2 area of the GoM is approximately 30 times larger (with considerably more complicated bathymetry and dynamics), yet the model for the region included only three additional input wave buoy locations, and was trained using 9 and 19 years less data for the temporal and spatial model components, respectively. As such, this is therefore an example of extreme sparsity and the performance of the machine learning model considered particularly impressive in this context.

Increasing the quantity and quality of information available for model training, as well as the inclusion of further model forcing locations in more optimal positions for capturing the evolution of conditions within both deep and shallow water, are likely key to improving performance. Extending the amount of usable (concurrent) observations for training the temporal model may be achieved by more advanced methods of gap filling that exploit temporal and spatial dependencies rooted in the wave buoy data, such as a low rank tensor completion with truncated nuclear norm algorithm (Chen et al., 2022), with the previous validation of such approaches demonstrating these are capable of restoring datasets missing up to 20% of their values with a coefficient of determination (r^2) exceeding 0.9. Extending the period of training for the spatial model may be achieved by using the full extent of the available hindcast (i.e. 1 January 1980 to 31 December 2022), with an additional advantage potentially able to be leveraged by using these data at their native spherical multiple cell resolution. It is noted that while these products generally perform very well, the temporal comparison of the output of the hindcast with that of the independent observations from the wave buoy in the vicinity of the Louisiana Offshore Oil Port shows that, particularly in coastal areas, differences (of up to ~0.5m at times) can still exist and these errors will then be transferred to the (learned) spatial surrogate. In addition to maximizing the exploitation of existing data, both satellites and ships can also provide additional methods of expanding the observation network via remote sensing and in-situ measurements. For example, by processing wave-induced motions from vessels for real-time sea state estimation at their locations, Mounet et al (2023) recently demonstrated these can be used as an input similarly to traditional observations. Integration of such response-based estimates have shown promising potential to supplement (or even replace) wave buoy records when used in nowcast-only mode (Mounet et al., 2023), increasing the accuracy and precision of the output spatial wave data over a large extent of the computational domain, providing the basis for possible further application in full spatiotemporal mode (especially in locations where more traditional observations platforms are scarce).

For reasons of scope and simplicity, the GoM configuration was constrained to the short-range prediction of wave conditions, particularly H_s, up to 12 hours ahead (recognizing that M_{Dir}, T_m and T_p are also output by the spatiotemporal model). However, the structure and flexibility of the framework suggests its potential further utility in related metocean challenges, such as in enabling enhanced prediction of complex LC/LCE dynamics within the region, by adopting the same coupled LSTM-RNN/RF machine learning approach. While this is still at an early stage of refinement, relevant work by Wang et al. (2019) and Muhamed-Ali et al. (2021) describing the development of an LSTM-RNN based model for the prediction of sea surface height and velocity structures in the same region have shown promising potential, hinting at the possible opportunities and benefits emerging from this type of approach. Indeed, within the GoM, the present authors see the application of the framework to the prediction of regional circulation patterns (rather than simply regional wave conditions) as being most significant in achieving real advantage for the industry from improved forecasts within the basin, owing jointly to the adverse impacts of LC/LCE presence on operational downtime (Steele et al., 2023), as well as the additional attainable accuracy resulting from the ability to directly use the full extent of the available real-time in-situ data to drive the model (Chen et al. 2021) – with the inclusion of new high-frequency radar-based surface current velocity observations across the Yucatan Channel likely of particular importance in this respect. For example, from an assessment of the spatiotemporal variability of the Yucatan Current (YC) using such data, DiMarco et al. (2023) established a dependence of the LC configuration on the location of the YC speed core, finding that the LC follows a more direct (retracted) inflow-outflow path when the speed core is near the channel midpoint and the LC extends further into the basin when the speed core moves toward the western side. Since rapid transitions in YC transports (e.g. associated with LCE shedding events) are not always well captured by traditional physics-based methods, the inclusion of these measurements is arguably able to be achieved most effectively and efficiently using data-driven approaches; crucially requiring no substantial alteration to the existing framework.

Future applications will focus on extending forecast horizon, the integration of additional observations and training datasets for the prediction of a wider range of metocean parameters (particularly currents), as well as the further development of the concept into a fully relocatable capability, able to be readily deployed, tested and applied to support improved decision-making in other regions.

Acknowledgements

Initial development of the framework presented was funded by the EPSRC-funded SuperGen Offshore Renewable Energy Hub (Grant EP/S000747/1) under the Machine Learning for Low-Cost Offshore Modelling (MaLCOM) project. A. C. Pillai acknowledges support from the Royal Academy of Engineering under the Research Fellowship scheme (Award RF\202021\20\175). The authors would like to thank Jessica Standen (Met Office) for her assistance, discussion and feedback.

References

Bi K, Xie L, Zhang H, Chen X, Gu X & Tian Q (2022) Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, **619**, 533-538. doi: https://doi.org/10.1038/s41586-023-06185-3.

Breiman L (2001) Random forests. Machine learning. 45, 5-32. doi: https://doi.org/10.1023/A:1010933404324.

Chen J, Pillai AC, Johanning L & Ashton IGC (2021) Using machine learning to derive spatial wave data: a case study for a marine renewable energy site. *Environmental Modelling & Software*, **142**, 105066. doi: https://doi.org/10.1016/j.envsoft.2021.105066.

Chen J, IGC Ashton & Pillai AC (2022) Wave record gap-filling using a low-rank tensor completion model. *ASME 41st International Conference on Ocean, Offshore and Arctic Engineering (OMAE2022)*, Hamberg, Germany, ASME, OMAE2022-79897. doi: https://doi.org/10.1115/OMAE2022-79897.

Chen J, Ashton IGC, Steele ECC & Pillai AC (2023) A real-time spatiotemporal machine learning framework for the prediction of nearshore wave conditions. *Artificial Intelligence for the Earth Systems*, **2**, e220033. doi: https://doi.org/10.1175/AIES-D-22-0033.1.

DiMarco SF, Glenn S, Smith M, Ramos R, Knap AH, Jimenez RM, Salas de Leon D & Tereza VKC (2023) Surface current velocity observations of the Yucatan Channel using high-frequency radar. *OCEANS 2023 – MTS/IEEE U.S. Gulf Coast*, Biloxi, MS, USA, 25-28 September 2023. doi: https://doi.org/10.23919/OCEANS52994.2023.10337389.

Du S, Li T, Yang Y & Horng S-J (2020) Multivariate time series forecasting via attention-based encoder-decoder framework. *Neurocomputing*, 388, 269-279. doi: https://doi.org/10.1016/j.neucom.2019.12.118.

Fan S, Xiao N & Dong S (2020) A novel model to predict significant wave height based on long short-term memory network. *Ocean Engineering*, **205**, 107298. doi: https://doi.org/10.1016/j.oceaneng.2020.107298.

Hersbach H, Bell B, Berrisford P, Hirahara S, Horanyi A, Munoz-Sabater J, Nicolas J, Peubey C, Radu R, Schepers D, Simmons A, Soci C, Abdalla S, Abellan X, Balsamo G, Bechtold P, Biavati G, Bidlot J, Bonavita M, DeChiara G, Dahlgren P, Dee D, Diamantakis M, Dragani R, Flemming J, Forbes R, Fuentes M, Geer A, Haimberger L, Healy S, Hogan RJ, Holm E, Janiskova M, Keeley S,Laloyaux P, Lopez P, Lupu C, Radnoti G, deRosnay P, Rozum I, Vamborg F, Villaume S & Thepaut J-N (2020) The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, **146** (730), 1999-2049. doi: https://doi.org/10.1002/qj.3803.

James SC, Zhang Y & O'Donncha F (2018) A machine learning framework to forecast wave conditions. *Coastal Engineering*, **137**, 1-10. doi: https://doi.org/10.1016/j.coastaleng.2018.03.004.

Keisler R (2022) Forecasting global weather with graph neural networks. *arXiv*, 2202.07575. doi: https://doi.org/10.48550/arXiv.2202.07575.

Lam R, Sanchez-Gonzalez A, Willson M, Wirnsberger P, Fortunato M, Alet F, Ravuri S, Ewalds T, Eaton-Rosen Z, Hu W, Merose A, Hoyer S, Holland G, Vinyals O, Stott J, Pritzel A, Mohamed S & Battaglia P (2023) Learning skillful medium-range global weather forecasting. *Science*, **382** (6677), 1416-1421. doi: https://doi.org/10.1126/science.adi2336.

Li J-G (2012) Propagation of ocean surface wavevs on a spherical multiple-cell grid. *Journal of Computational Physics*, **231**, 1536-1555. doi: https://doi.org/10.1016/j.jcp.2012.08.007.

Luong M-T, Pham H & Manning CD (2015) Effective approaches to attention-based neural machine translation. *arXiv*, 1508.04025v5. doi: http://arxiv.org/abs/1508.04025.

Meyer H, Reudenbach C, Hengl T, Katurji M & Nauss T (2018) Improving performance of spatio-temporal machine learning models using forward feature selection and target-oriented validation. *Environmental Modelling & Software*, **101**, 1-9. doi: https://doi.org/10.1016/j.envsoft.2017.12.001.

Molnar C (2020) Interpretable machine learning: a guide for making black box models explainable. https://christophm.github.io/interpretable-ml-book/tree.html#interpretation-2.

Mounet REG, Chen J, Nielsen UD, Brodtkorb AH, Pillai AC, Ashton IGC & Steele ECC (2023) *Ocean Engineering*, **281**, 114892. doi: https://doi.org/10.1016/j.oceaneng.2023.114892.

Muhamed-Ali A, Zhuang H, VanZwieten J, Ibrahim AK & Cherubin LM (2021) A deep learning model for forecasting velocity structures of the Loop Current System in the Gulf of Mexico. *Forecasting*, **3** (4), 934-953. doi: https://doi.org/10.3390/forecast3040056.

Pathak J, Subramanian S, Harrington P, Raja S, Chattopadhyay A, Mardani M, Kurth T, Hall D, Li Z, Azizzadenesheli K, Hassanzadeh P, Kashinath K & Anandkumar A (2022) Fourcastnet: A global data-driven high-resolution weather model using adaptive Fourier neural operators. *arXiv*, 2202.11214. doi: https://doi.org/10.48550/arXiv.2202.11214.

Pirhooshyaran M, Scheinberg K & Snyder LV (2020) Feature engineering and forecasting via derivative-free optimization and ensemble of sequence-to-sequence networks with applications in renewable energy. *Energy*, **196**, 117136. Doi: https://doi.org/10.1016/j.energy.2020.117136.

Pirhooshyaran M & Snyder LV (2020) Forecasting, hindcasting and feature selection of ocean waves via recurrent and sequence-to-sequence networks. *Ocean Engineering*, **207**, 107424. doi: https://doi.org/10.1016/j.oceaneng.2020.107424.

Saulter A, Bunney C & Li J-G (2016) Application of a refined grid global model for operational wave forecasting. Met Office Research Technical Report 614, 46. doi: https://doi.org/10.13140/RG.2.2.22242.17600.

Schultz MG, Betancourt C, Gong B, Kleinert F, Langguth M, Leufen LH, Mozaffari A & Stadtler (2021) Can deep learning beat numerical weather prediction. *Philosophical Transactions of the Royal Society*, **A379**, 20200097. Doi: https://doi.org/10.1098/rsta.2020.0097.

Sinha A & Abernathey R (2021) Estimating ocean surface currents with machine learning. *Frontiers in Marine Science*, **8**, 672477. doi: https://doi.org/10.3389/fmars.2021.672477.

Song T, Pang C, Hou B, Xu G, Xue J, Sun H & Meng F (2023) A review of artificial intelligence in marine science. *Frontiers in Earth Science*, **11**, 1090185. doi: https://doi.org/10.3389/feart.2023.1090185.

Sonnewold M, Lguensat R, Jones DC, Dueben PD, Brajard J & Balaji (2021) Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environmental Research Letters*, **16**, 073008. doi: https://doi.org/10.1088/1748-9326/ac0eb0.

Steele ECC, Jaramillo S, Neal R, Storie J & Zhang M (2023) Prediction of Loop Current and Eddy Regimes in the Gulf of Mexico. *Offshore Technology Conference*, Houston, TX, USA, 1-4 May 2023. doi: https://doi.org/10.4043/32615-MS.

Sutskever I, Vinyals O & Le QV (2014) Sequence to sequence learning with neural networks. *Proceedings of the 27th International Conference on Neural Information Processing Systems*, Volume 2 (NIPS'14), MIT Press, Cambridge, MA, USA, 3104-3112. https://proceedings.neurips.cc/paper_files/paper/2014/file/a14ac55a4f27472c5d894ec1c3c743d2-Paper.pdf.

Valiente NG, Saulter A, Gomez B, Bunney C, Li J-G, Palmer T & Pequignet C (2023) The Met Office operational wave forecasting system: the evolution of the regional and global models. *Geoscientific Model Development*, **16**, 2515-2538. doi: https://doi.org/10.5194/gmd-16-2515-2023.

Wang JL, Zhuang H, Cherubin LM, Ibrahim AK & Muhamed-Ali A (2019) Medium-term forecasting of Loop Current Eddy Cameron and Eddy Darwin formation in the Gulf of Mexico with a divide-and-conquer machine learning approach. *Journal of Geophysical Research Oceans*, **124** (8), 5586-5606. doi: https://doi.org/10.1029/2019JC015172.

Appendix 1

While the RF-based spatial model utilizes all of the available input data to define the mapping between the individual wave buoy locations and the rest of the model domain concurrently, from the entire sensor network rather than just that of the nearest site, it is an obvious fact that extreme sparsity nonetheless adversely affects its ability to reconstruct the spatial field – and this, in turn, is somewhat dependent on distance (i.e. locations further away will be more impacted). Therefore, as an indication of areas where worse performance may be expected, Figure 5 presents a map of the mean coefficient of determination (r^2) at each of the grid points for the period 1 January 2021 to 31 December 2022, associated with its most correlated sensor in the network. Increasing the density of input observations, as well as optimizing the locations of their collection relative to the configuration and size of the model domain, is deemed important for increasing performance compared to traditional physics-based methods.

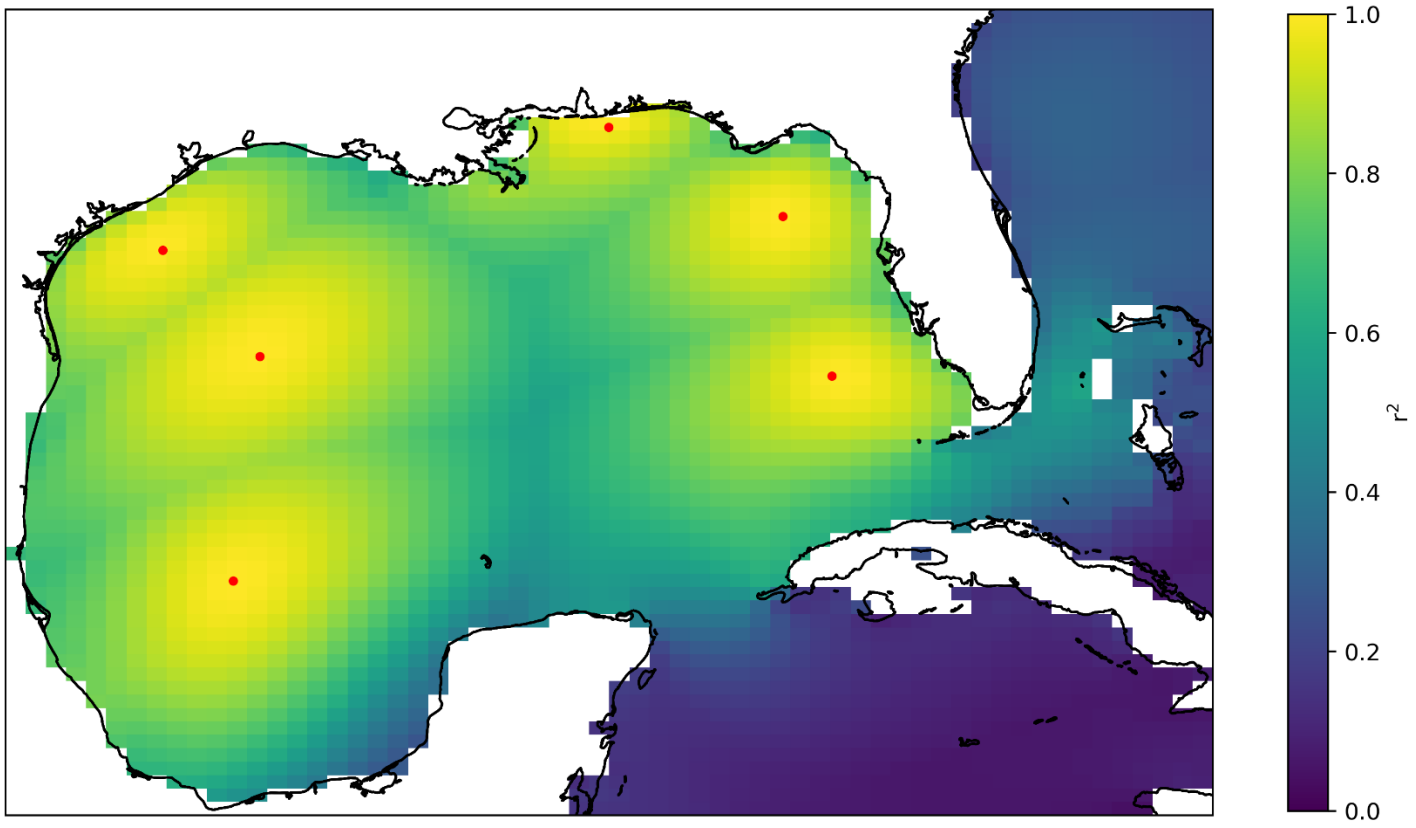

Figure 5: A map of the mean coefficient of determination (r^2) at each of the grid points for the period 1 January 2021 to 31 December 2022, associated with its highest correlated sensor in the network.

Appendix 2

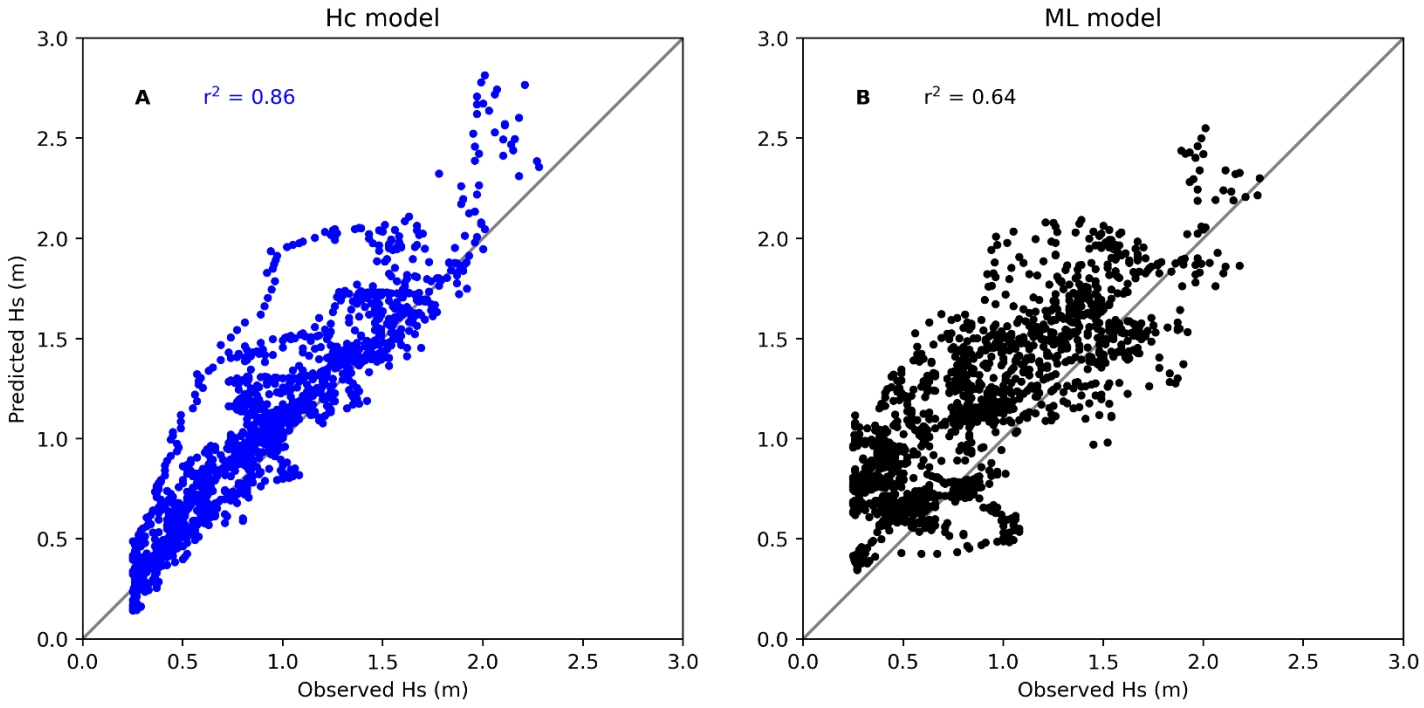

Figure 6: Scatterplots of physics-based hindcast model (left hand column) and machine learning model (right hand column) performance, relative to the buoy observations, at a lead time of T+0. The diagonal gray line shows an (ideal) 1:1 relationship. Note that this is simply an alternative presentation of the same comparison presented as a timeseries in Figure 4A.