Offshore Technology Conference

OTC-35733-MS

A Spatiotemporal Machine Learning Framework for the Prediction of Metocean Conditions in the Gulf of Mexico: Application to Loop Current and Loop Current Eddy Forecasting

E.C.C. Steele[1], M.C.R. Juniper[1], A.C. Pillai[2], I.G.C. Ashton[2], J. Chen[3], S. Jaramillo[4] & L. Zarate[5]

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

Abstract

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. By capitalizing on a vast amount of available information – and capturing their inherent patterns that are not represented explicitly – such machine learning-based techniques have the potential to increase forecast accuracy, augmenting traditional physics-based equivalents. Here, we adapt and apply a promising machine learning framework – originally proposed by the present authors for regional prediction of ocean waves – to the operational forecasting of the Loop Current and Loop Current Eddies (LC/LCEs) in the Gulf of Mexico (GoM). 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 observations, that is then combined with a random forest based spatial nowcasting model, trained on high-resolution regional reanalysis data, to develop a complete spatiotemporal prediction for the basin. Since machine learning approaches are typically physics-agnostic, an identical framework to that developed for the prediction of ocean waves can be used for the prediction of surface currents, with the only difference being the training datasets to which this is exposed. This is illustrated here using a period of three months of training data from October 2022 to December 2022, with the model driven by only three observation sites in the northern GoM. As such, it is unrealistic to expect performance for an unseen week in January 2023 to be equivalent to smaller/simpler domains with a more favorable quantity, quality and coverage/distribution of input observations but, despite these severe constraints, the ability of the model to forecast a plausible structure of the LC/LCE system is nonetheless impressive. The architecture of the MaLCOM framework allows for easy interrogation of the temporal and spatial behavior of the model which allows us to better unpick and explain its characteristics – thus providing a path to inform further enhancements. While still at an early stage of refinement, the extension of the framework from waves to currents demonstrates encouraging potential for a fundamentally different approach to the way that metocean data in general, and LC/LCE forecasts in particular, can be generated and used by the offshore energy sector, by directly leveraging sparse sensor networks as the basis for these predictions (further extending the value of the observations, when collected with this additional purpose in mind). Provided a suitable coverage, quality and quantity of observations are available, the advent of these very low cost, data-driven predictions – able to be run on-demand, in-house, using standard laptop or desktop computers – herald new opportunities for improving real-time decision-making to support offshore planning and workability.

Introduction

Accurate forecasting of the frontal position and intensity of the Loop Current and Loop Current Eddies (LC/LCEs) in the Gulf of Mexico (GoM) – with a level of confidence and lead time to permit efficient industry decision-making – arguably represents the pinnacle of oceanographic research efforts in the region (National Academies of Sciences, Engineering & Medicine, 2018; Steele et al., 2023). Such features are typically associated with intense currents of up to 2 ms^{-1} (Oey et al., 2005a), with the prediction of their encroachment on offshore energy assets of critical importance to operators because of the adverse impacts on effective planning and workability (Sharma et al., 2016). However, despite decades of study (Sturges & Lugo-Fernandez, 2005; National Academies of Sciences, Engineering & Medicine, 2018; Morey et al., 2024), involving extensive monitoring (e.g. Schmitz, 2005) and modelling (e.g. Oey et al., 2005b)

initiatives, uncertainties about fundamental controls on the extent of penetration of the LC and the processes responsible for the LCE shedding persist; meaning that skill in accurately forecasting their asset-level location, duration and evolution remains elusive (Steele et al., 2023).

While predictive skill still remains somewhat limited, an upside of having such a vast amount of existing information available – including in-situ and remote sensing measurements, as well as numerical reanalyses and hindcasts – are the opportunities that this affords for the application of new machine learning techniques; the emergence of which is presently defining a new era of data-driven weather prediction. In contrast to traditional physics-based numerical weather prediction (NWP) methods that seek to directly solve the governing equations of fluid dynamics, machine learning-based weather prediction (MLWP) methods seek to identify and utilize the inherent patterns within data that are not represented explicitly. Advances in deep learning, such as the development of multi-layer perceptrons (MLPs), convolutional neural networks (CNNs), and recurrent neural networks (RNNs), combined with improved computational accessibility, has greatly enhanced MLWP capabilities (LeCun et al., 2015; Sonnewald et al., 2021). Indeed, recent years have seen a proliferation of MLWP models that demonstrate comparable performance with NWP equivalents for a subset of variables, although efforts have largely focused on global atmospheric forecasting (e.g. Keisler, 2022; Pathak et al., 2022; Bi et al., 2023; Lam et al., 2023; Bodnar et al., 2024; Lang et al., 2024) and, with the exception of Vaughan et al. (2024) and Alexe et al. (2024), still must be initialized from a physics-based (re)analysis. At the same time, it is recognized there is growing interest in the complementary development of regional MLWP models of metocean variables such as sea surface heights (e.g. Chattopadhyay et al., 2024), surface currents (e.g. Muhamed-Ali et al., 2021; Sinha & Abernathey, 2021) and waves (e.g. James et al., 2018; Fan et al., 2020), however, the potential for providing a standard, scalable and applications-based framework – that is both region and parameter agnostic has not yet been realized (Song et al., 2023, Steele et al., 2024).

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 regional prediction of ocean waves (Chen et al., 2021; Chen et al., 2023) – to the operational forecasting of the LC and LCEs in the GoM. The approach consists of using an attention-based long short-term memory recurrent neural network (LSTM-RNN) to learn the temporal patterns from a network of available observations, combined with a random forest (RF)-based spatial nowcasting model trained on reanalysis data, to develop a complete spatiotemporal prediction for the entire region (Steele et al., 2024). Since this machine learning approach is inherently physics-agnostic, an identical framework to that developed for the prediction of ocean waves can be used for the prediction of surface currents, with the only difference being the training datasets to which this is exposed. This is illustrated using a period of three months of training data from October 2022 to December 2022, with the model driven by only three observation sites in the northern GoM. As such, it is unrealistic to expect performance for an unseen week in January 2023 to be equivalent to smaller/simpler domains with a more favorable quantity, quality and coverage / distribution of input observations but, despite these severe constraints, the ability of the model to forecast a plausible structure of the LC/LCE system is nonetheless impressive. The architecture of the MaLCOM framework allows for easy interrogation of the temporal and spatial behavior of the model which allows us to better unpick and explain its characteristics – thus providing a path to inform further enhancements. While still at an early stage of refinement, the extension of the MaLCOM framework from waves to currents demonstrates encouraging potential for a fundamentally different approach to the way that metocean data in general, and LC/LCE forecasts in particular, can be generated and used by the offshore energy sector, by directly leveraging sparse sensor networks as the basis for these predictions (further extending the value of the observations, when collected with this additional purpose in mind). Provided a suitable coverage, quality and quantity of observations are available, the advent of these very low cost, data-driven predictions – able to be run on-demand, in-house, using standard laptop or desktop computers – herald new opportunities for improving real-time decision-making to support offshore planning and workability.

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 being extended, while the ‘Assessment’ section presents an evaluation of its potential for short-range forecast purposes. The final section presents a discussion of the concept, as well as outlining opportunities for further work.

Models & Methods

The complete framework consists of combining a site-specific LSTM-RNN temporal forecasting model (trained and initialized from observations; Chen et al., 2023) with a RF spatial nowcasting model (trained from a regional physics-based reanalysis; Chen et al., 2021). For the GoM basin, this configuration is identical to that described by Steele et al. (2024), with the exception being the observations and reanalysis datasets used.

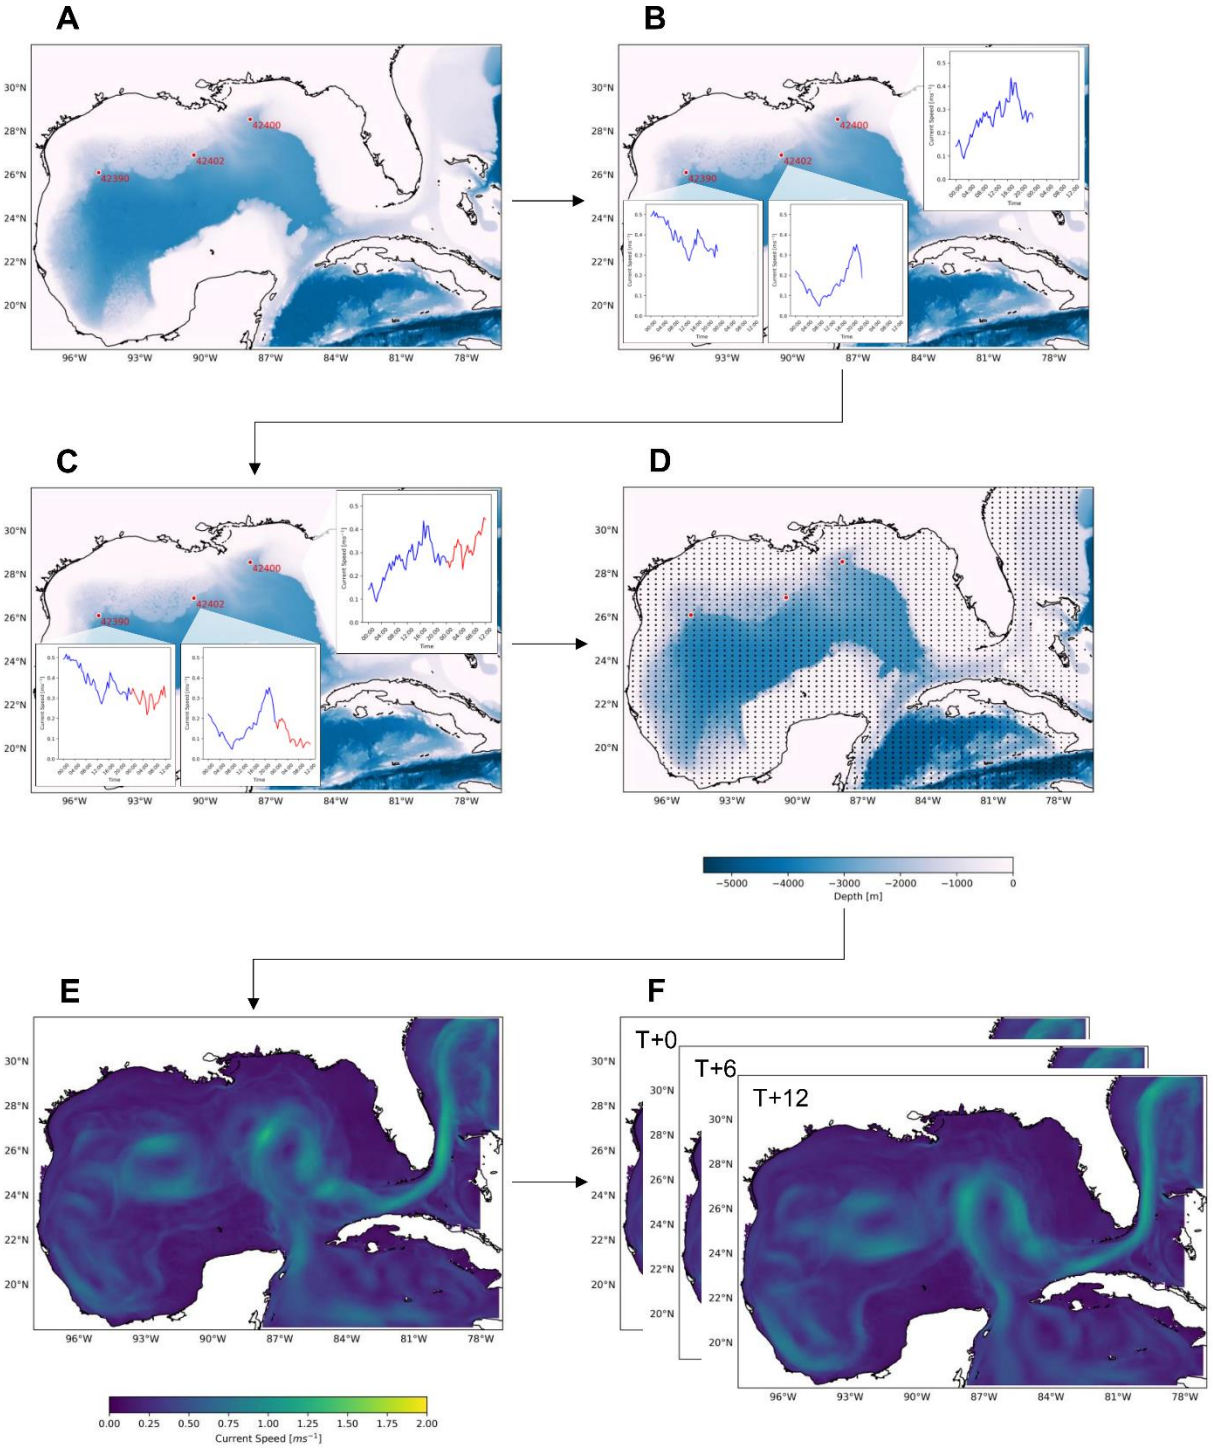

Figure 1: Overview of the MaLCOM framework. Here, a network of three observations locations (red dots; panel A) were used to learn the temporal patterns in current velocities based on the previous 24 hours of measurements (panel B), enabling a prediction at each of these three sites to be made up to 12 hours ahead (panel C), and the velocities associated with these locations can then be mapped across the region of interest (panel D). By providing the latest 24 hours of in-situ observations data (i.e. 48-half-hourly time steps) at each of the platform locations, the framework will therefore return the full spatial prediction of the surface currents across the domain (panel E), with a similar mapping able to be applied iteratively to the output from each time step of the temporal model, to give a full spatiotemporal forecast (panel F).

A simplified overview of this process is illustrated in Figure 1, and only a summary is provided here. Briefly, having first specified the regional domain of interest, several suitable platform locations must be identified from which coincident time-series observations are available (Figure 1A). The three platform locations shown here – all equipped with Acoustic Doppler Current Profiler (ADCP) instruments – were selected as they provided relatively complete, continuous measurements of near-surface current

velocities with an hourly or higher sampling frequency (shown by the blue lines in Figure 1B). Assuming a reasonable quantity and quality of observations are available, a LSTM-RNN model can then be trained on these data, in order to learn the temporal patterns using the network of available observations – thereby enabling timeseries prediction of these parameters to be made at each of the individual locations up to 12 hours ahead (shown by the red lines in Figure 1C), based on the previous 24 hours of measurements (shown by the blue lines in Figure 1C). The velocities associated with predictions at the ADCP locations can then be mapped across the region of interest (Figure 1D) using a RF model that previously learned these relationships from an existing physics-based ocean reanalysis. By providing the latest 24 hours of in-situ observations data (i.e. 48-half-hourly time steps) at each of the platform locations, the framework will therefore return a spatial prediction of the surface currents across the domain (Figure 1E), with a similar mapping able to be applied iteratively to the output from each time step of the temporal model, to give a full spatiotemporal forecast (Figure 1F) up to 12 hours ahead (i.e. 24-half-hourly time steps) or, by tuning with a longer lookback period and/or time step, beyond.

For the present study, the temporal model was trained using three months of ADCP data from each of the three input ADCP locations (between 1 October 2022 and 31 December 2022) and the spatial model was trained using an identical three months of data from the 00Z and 12Z fields of the HYCOM-TSIS 1/25 Gulf of Mexico (GoMb0.04) reanalysis. While certainly limiting, these training constraints were imposed by the quantity and quality of available/coincident observations, as well as the very high resolution of the hindcast increasing the size of some of the resultant MaLCOM data structures in memory. For the training of the temporal model, the raw, publicly-available ADCP data at each site was clustered by depth, using a one-dimensional Kernel Density Estimation (KDE)-based method (Appendix 1) to ensure selection of data only from within the bin nearest the surface (~55-65m), with no further processing of these measurements conducted (e.g. low pass filtering of velocities, or rescaling from sub-surface to surface values via the incorporation of information about the shear/vertical profile). For the training of the spatial model, the reanalysis (GoMb0.04), used to learn the mapping between the input ADCP locations and the rest of the model domain, was based on the Hybrid Ocean Coordinate Model (HYCOM; Chassignet et al., 2007) and uses the Tendral Statistical Interpolation System (TSIS; Srinivasan et al., 2022) package for the assimilation of satellite-derived sea level anomaly (SLA) and sea surface temperature (SST) measurements, as well as in-situ vertical temperature and salinity profiles. A recent process-orientated assessment of the product in the northern GoM was conducted by Ivanov et al. (2024), showing good qualitative and satisfactory quantitative agreement between the observations and the reanalysis. For simplicity of initial testing, both LSTM-RNN and RF model components were presented with only the current speeds derived from the northwards and eastwards velocities from the respective ADCP or reanalysis datasets.

Assessment

An assessment of the spatiotemporal machine learning approach is conducted for the short-range prediction of LC/LCEs up to 12 hours ahead in the GoM using 5 days of unseen data from the start of January 2023.

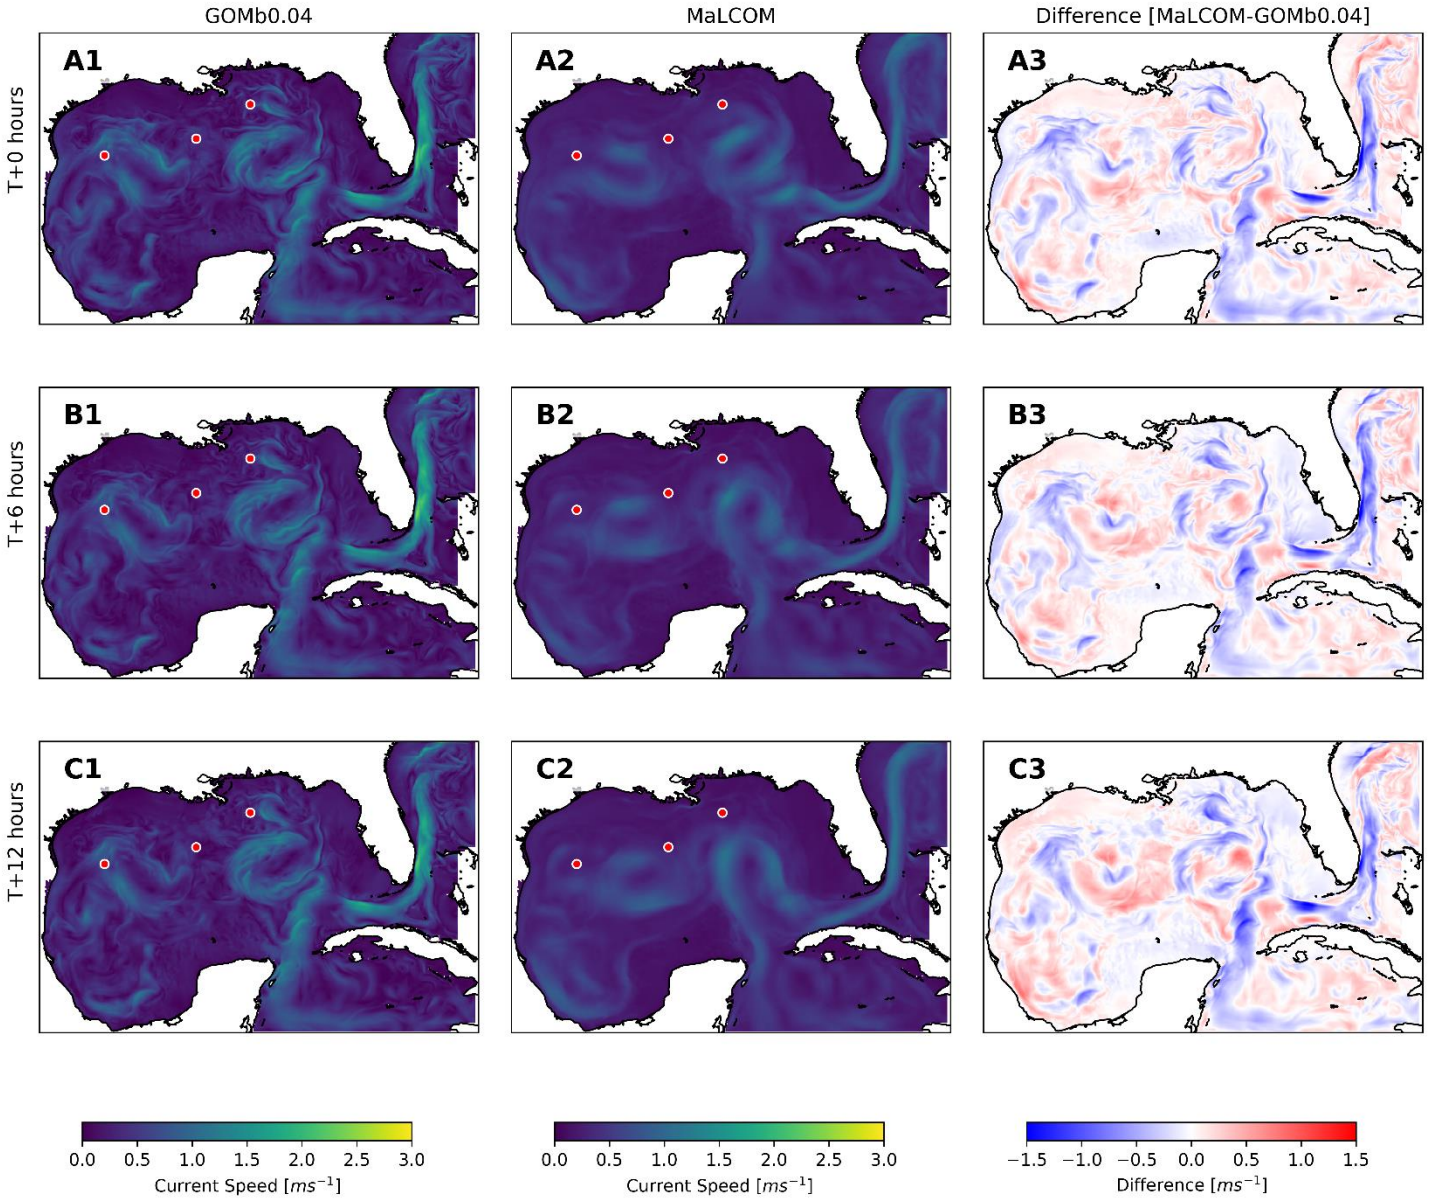

Figure 2: Example comparison of the surface current output of the physics-based reanalysis ('GOMb0.04') model with that of the machine learning ('MaLCOM') model, relative to the 00Z run on 3 January 2023. The current speed from the reanalysis model is shown in the left-hand column, the current speed obtained from the machine learning model is shown in the middle column, and the difference between the two (calculated as MaLCOM-GOMb0.04) 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 1, the overlaid markers indicate the three input ADCP locations (red dots).

Figure 2 presents an example comparison of the surface current output of the physics-based reanalysis ('GOMb0.04') model with that of the machine learning ('MaLCOM') model, relative to the 00Z run on 3 January 2023. Here, the current speed from the reanalysis model is shown in the left-hand column, the current speed obtained from the machine learning model is shown in the middle column, and the difference between the two (calculated as MaLCOM-GOMb0.04) 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. Encouragingly, despite the extreme sparsity of the input ADCP locations (red dots), it is seen that the machine learning model can forecast a plausible structure of the LC/LCEs from just three sites in the northern GoM. Here, it is seen there are two LCEs present in the eastern and western GoM respectively, with the latter one tentatively attached to the LC itself. The difference in predicted intensity of these current fields is a likely consequence of the constraints of the training duration and number of input features, combined with the discrepancy in velocity between the observations and the reanalysis, associated with the difference between the depth the ADCP measurements were collected (~55-65 m) and the surface level of the numerical reanalysis (~0 m). The machine learning model exhibits stable performance with only a very small increase in the spatially-averaged mean absolute error (MAE) from 0.16 ms^{-1} to 0.18 ms^{-1} and root mean squared error (RMSE) from 0.05 ms^{-1} to 0.06 ms^{-1} with lead time. However, consistent with all data-driven approaches,

but particularly under the conditions of extreme sparsity described here, this is likely due to being close to the mean of the training data probability distribution function, leading to progressively overly-smoothed fields with the MAE and RMSE increasingly close to climatology with lead time (Bonavita et al., 2024).

While Figure 2 illustrates an impressive potential for the machine learning model to predict surface currents from only a small quantity of ultra-sparse observations, a longer period for both training and testing the model (across a variety of LC/LCE states) is needed to properly characterize its performance. This evaluation is clearly beyond the scope of the present study, but it is possible to interrogate the individual temporal and spatial components of the MaLCOM framework to explain its behavior.

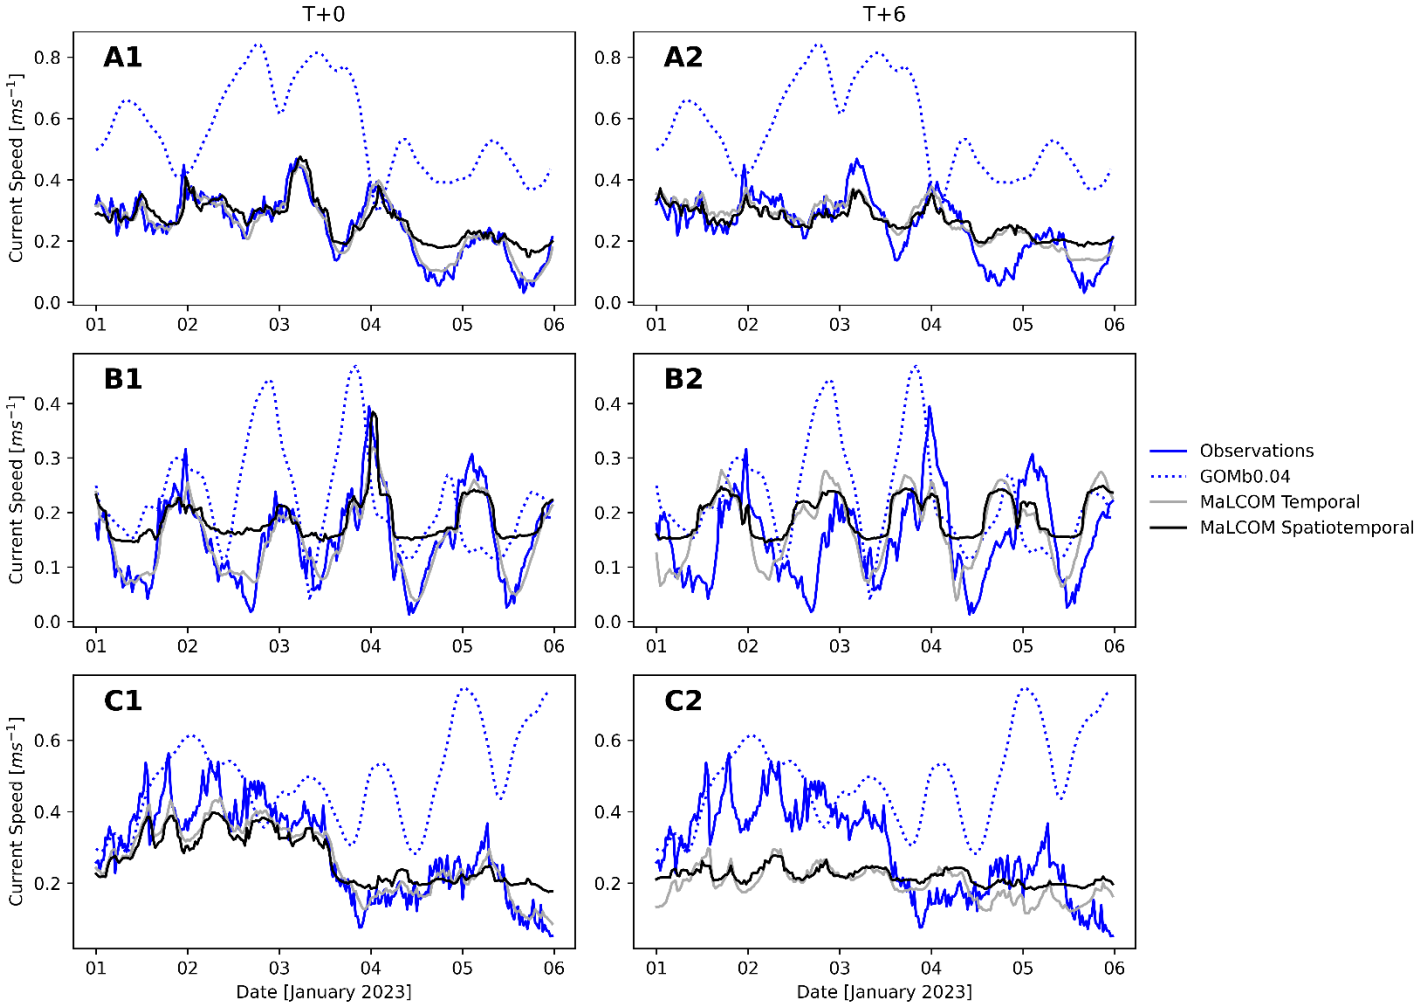

Figure 3: Example comparison of the output of the MaLCOM temporal prediction (gray line) and the output of the full MaLCOM spatiotemporal prediction (black line), with that of the observations (solid blue line) and the physics-based reanalysis ('GOMb0.04') model (dashed blue line) for each of the three input ADCP locations (top row: site 42390; middle row: site 42402; bottom row: site 42400), at a lead time of T+0 (left-hand column) and T+6 (right-hand column) hours ahead.

Figure 3 presents an example comparison of the output of the MaLCOM temporal prediction (gray line) and the output of the full MaLCOM spatiotemporal prediction (black line), with that of the observations (solid blue line) and the physics-based reanalysis ('GOMb0.04') model (dashed blue line) for each of the three input ADCP locations (top to bottom rows, from west to east across the GoM, respectively), at a lead time of T+0 (left-hand column) and T+6 (right-hand column) hours ahead. These reveal that initially (i.e. T+0) the temporal model captures the pattern of the observations well, with performance better in the western GoM than in the eastern GoM (likely associated with these locations experiencing less variability within the period used for both training and testing). When the temporal model is combined with the spatial model to enable full spatiotemporal prediction, a slight degradation in performance is seen (likely associated with the discrepancy in velocity of up to ~0.3 ms^{-1} due to the difference between the depth the ADCP measurements were collected (~55-65 m) and the surface level of the numerical reanalysis (~0 m), which are independently used for training the two individual components). By T+6, the differences between the observations and the predictions increase, with a loss of amplitude across all locations, albeit particularly apparent in the eastern GoM. While this

damping is certainly compounded by the limited quantity of training data, as well as a possible influence arising from other sites within the network of available observations, it is possible that further performance gains could be achieved by extending the lookback period and/or time step to better capture the longer-term patterns of temporal variability; in addition to increasing the total number of input locations/features in more suitable positions for optimizing the performance of the combined system.

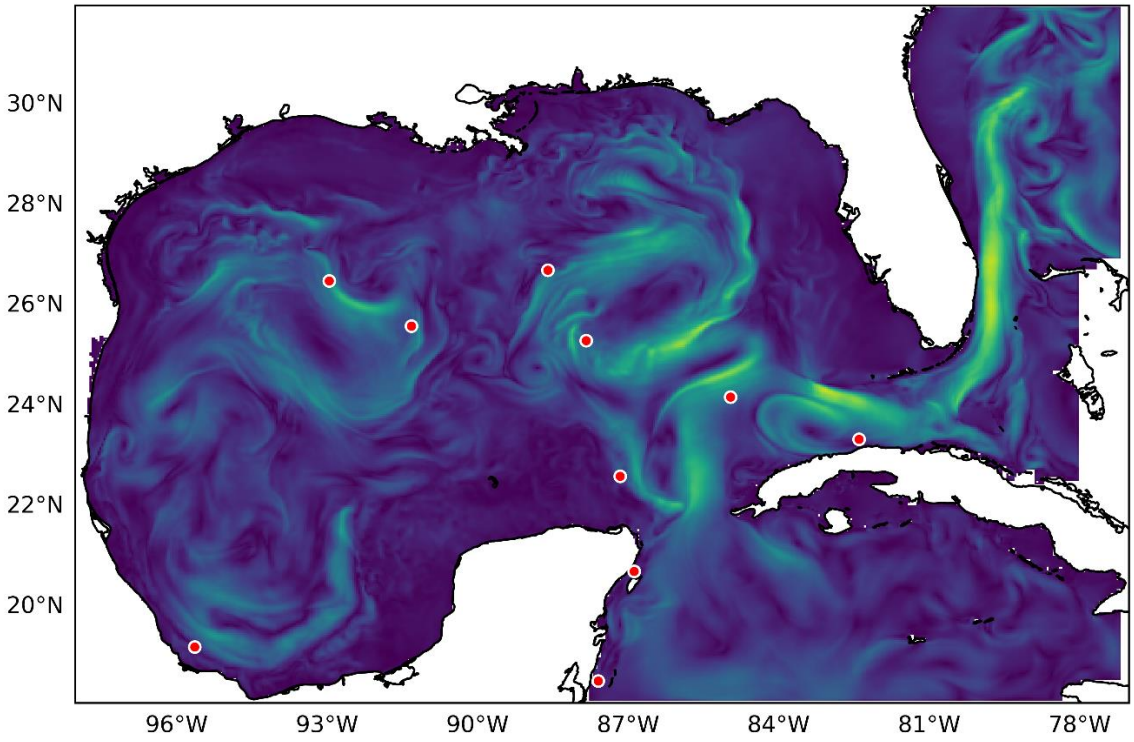

Figure 4: Example optimization of sensor number and placement, obtained by minimizing the reconstruction error within the three-month training period, with the aim of providing insight into areas with the most potential to improve the machine learning forecast accuracy itself. The optimized locations of 10 sensors within both the interior of the GoM itself as well as its upstream inflow region (red dots) are overlaid onto an example snapshot of the surface current output from the reanalysis (GOMb0.04) model at the end of the three-month training period.

Since increasing both the quantity and quality of information available – particularly via the inclusion of further input locations/features in more suitable positions for capturing the evolution of conditions – is so critical, Figure 4 presents an example optimization of sensor number and placement, obtained by minimizing the reconstruction error of the reanalysis model over the three-month training period (deSilva et al., 2021), using 10 sensors within both the interior of the GoM itself as well as its upstream inflow region, with the aim of providing insight into areas with the most potential to improve the machine learning forecast accuracy itself. Perhaps unsurprisingly, there are a greater number of observations required in the south-eastern part of the domain (consistent with complex dynamics associated with the inflow/outflow conditions in the Yucatan Channel and Florida Channel), with these stretching northwards and westwards (aligned with the direction of typical penetration of the LC and propagation of LCEs) – reaffirming that incorporating new measurements from these areas (notably absent within the present study) are essential.

Discussion

These results demonstrate the extension of the MaLCOM framework from waves to currents, enabling the prediction of a larger set of metocean conditions/processes than considered previously, using a parameter-agnostic approach capable of taking advantage of both rapidly updating observations and existing regional physics-based reanalysis (Steele et al., 2024). While opportunities for the further development of the capability for LC/LCE application are acknowledged, the benefit of such low-cost, data-driven predictions – enabling, for example, the issue of more frequent forecast updates at much higher resolution – offer exciting potential to augment traditional metocean prediction methods that have long been critical to supporting the safe and efficient operation of offshore infrastructure.

For the present study, only three months of data (from 1 October 2022 to 31 December 2022), from just a single forecast variable at three sites in the northern GoM, were used to train both the LSTM-RNN and the RF model components. Compared with the previous application of the framework for the prediction of wave conditions on the European North West Shelf (NWS; Chen et al.,

2023) and in the GoM (Steele et al. 2024), the constraints imposed here are considerably more extreme. Between 6 months and 21 years of data, including four forecast variables, were variously used across these prior cases; with three locations being used in the NWS waves application (whose domain is 30-times smaller in size) and six locations being used in the GoM waves application (whose positioning relative to relevant processes responsible for driving the spatiotemporal variability throughout the region was also arguably more suitable). In this context, the ability of the model to forecast a plausible structure of the LC/LCEs over an area 1,600,000 km^2 in size under such challenging conditions is therefore considered highly encouraging.

While it is inevitable that this performance is in part associated with the similarity of the unseen test conditions to that used for training (as a consequence of availability/resolution constraints), this is not inconsistent with adopting an approach where a series of machine learning models are pre-built for a variety of historical circulation types – or 'regimes' – that represent the large scale LC/LCE 'state' in the GoM; with the one conditioned for the closest matching regime to the prevailing classification subsequently applied in inference. An example set of 40 general-purpose LC/LCE regime definitions, suitable for both historical and real-time application, has previously been proposed by Steele et al. (2023), with it suggested these can be leveraged to enable the strategic optimization of machine learning based metocean forecasts in operation. In addition to accuracy, an auxiliary advantage of employing such a regime-based concept is that it might enable the size of the individual machine learning models themselves to be reduced – permitting a longer training period, without compromising performance – thus making this even more suitable for low-cost applications (including potentially also onboard the assets themselves).

By separating the spatial and temporal prediction elements, the structure of the MaLCOM framework promotes explainability of its behavior, enabling subsequent development effort to be focused on specific components, as well as the interaction of these in combination. Fundamentally, however, as with all machine learning techniques, increasing both the quantity and quality of information available, and the inclusion of further input locations in more optimal positions for capturing the evolution of conditions, are critical for improving performance (Steele et al., 2024). Indeed, the increased adoption of data-driven approaches places a renewed importance on the coordination and collection of continuous observation given their direct value in forecasting – recognizing appropriate/consistent quality control (often lacking in many publicly-available datasets) is also essential. Traditionally, the location of such in-situ velocity measurements have been limited to platform locations and/or in proximity to specific features (such as LC/LCE fronts), to best enable asset-based monitoring, but an alternative approach might be to look to target areas with the most potential to improve the machine learning forecast accuracy (and/or spatiotemporal reconstruction) itself. Examples may include the incorporation of the new high-frequency radar-based surface current observations across the Yucatan Channel (the position of the speed core of which has a known impact on downstream LC configuration; DiMarco et al., 2023); or the strategic use of drifting buoys and/or ships (that are routinely used to complement the daily frontal analysis of the behavior of the LC system; Storie et al., 2023). Incorporation of moving observations within the MaLCOM framework may be readily achieved by implementing a convolutional conditional neural process (Gordon et al., 2019), to map these to a set of optimized locations that can then be used to drive the model; with the benefit of facilitating the handling of dropouts in observations to increase temporal duration, and the integration of satellite remote sensing datasets to further increase spatial coverage/distribution. In addition, the neural process permits estimation of the associated uncertainty, thus providing a path to probabilistic predictions deemed to be of benefit in user decision-making, and potentially even a means of replacing the reanalysis used for learning the spatial relationships between the points in the domain exclusively using observations (e.g. Vaughan et al., 2024).

While still at an early stage of refinement, the extension of the MaLCOM framework from waves to currents described here demonstrates encouraging potential for its further development and tuning. In addition to conducting real-time trials to support live operations in the GoM, future work will focus on extending the forecast horizon via further tuning of the framework for currents, particularly involving the integration of moving observations and the incorporation of regime-based principles – with such approaches expected to augment traditional physics-based forecast models.

Acknowledgements

A.C. Pillai acknowledges support from the Royal Academy of Engineering under the Research Fellowship scheme (award number: RF\202021\20\175). I.G.C. Ashton acknowledges support from the Royal Academy of Engineering under Industrial Fellowship scheme (award number: IF-2425-19-AI155). E.C.C. Steele acknowledges the National Academies of Sciences, Engineering & Medicine for enabling his participation in the UGOS Annual Meeting 2024 where some of these ideas were discussed. The authors are grateful to Jessica Standen (Met Office) for her feedback on the final draft of the manuscript.

References

Alexe M, Boucher E, Lean P, Pinnington E, Laloyaux P, McNally A, Lang S, Chantry M, Burrows C, Chrust M, Pinault F, Villeneuve E, Bormann N & Healy (2024) GraphDOP: towards skilful data-driven medium-range weather forecasts learned and initialized directly from observations. arXiv, 2412.15687. doi: https://doi.org/10.48550/arXiv.2412.15687.

Bi K, Xie L, Zhang H, Chen X, Gu X & Tian Q (2023) Accurate medium-range global weather forecasting with 3D neural networks. Nature, 619, 533-538. doi: https://doi.org/10.1038/s41586-023-06185-3.

Bodnar C, Bruinsma WP, Lucic A, Stanley M, Vaughan A, Brandstetter J, Garvan P, Riechert M, Weyn JA, Dong H, Gupta JK, Thambiratnam K, Archibald AT, Wu C-C, Heider E, Welling M, Turner RE & Perdikaris P (2024) A foundation model for the Earth system. arXiv, 2405.13063. doi: https://doi.org/10.48550/arXiv.2405.13063.

Bonavita M (2024) On some limitations of current machine learning weather prediction models. Geophysical Research Letters, 51(12), e2023GL107377, doi: https://doi.org/10.1029/2023GL107377.

Chassignet EP, Hurlburt HE, Smedstad OM, Halliwell GR, Hogan PJ, Wallcraft AJ, Baraille R & Bleck R (2007) The HYCOM (Hybrid Coordinate Ocean Model) data assimilative system. Journal of Marine Systems 65 (1-4), 60-83. doi: https://doi.org/10.1016/j.jmarsys.2005.09.016.

Chattopadhyay A, Gray M, Wu T, Lowe AB & He R (2024) OceanNet: a principled neural operator-based digital twin for regional oceans. Scientific Reports 14, 21181. doi: https://doi.org/10.1038/s41598-024-72145-0.

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, 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.

deSilva BM, Manohar K, Clark E, Brunton BW, Brunton SL & Kutz JN (2021) PySensors: a python package for sparse sensor placement. arXiv, 2102.13476. doi: https://doi.org/10.48550/arXiv.2102.13476.

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.

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.

Gordon J, Bruinsma WP, Foong AYK, Requeima J, Dubois Y & Turner RE (2019) Convolutional conditional neural processes. arXiv, 1910.13556. doi: https://doi.org/10.48550/arXiv.1910.13556.

Ivanov L, Arena R, Bozec A, Chassignet E, Longridge S, Ramos R, Srinivasan A & Iskandarani M (2024) Process-oriented validation of HYCOM-TSIS reanalysis runs for the northern Gulf of Mexico. Offshore Technology Conference, Houston, TX, USA, 6-9 May 2024. doi: https://doi.org/10.4043/35402-MS.

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.

Lang S, Alexe M, Chantry M, Dramsch J, Pinault F, Raoult B, Clare MCA, Lessig C, Maier-Gerber M, Magnusson L, Bouallegue ZB, Nemesio AP, Dueben PD, Brown A, Pappenberger F & Rabier F (2024) AIFS – ECMWF's data-driven forecasting system. arXiv, 2406.01465. doi: https://doi.org/10.48550/arXiv.2406.01465.

LeCun Y, Bengio Y & Hinton G (2015) Deep Learning. Nature, 521, 436-444. doi: https://doi.org/10.1038/nature14539.

Morey SL, He R, Chassignet EP & Sheinbaum J (2024) Editorial: understanding and predicting the Gulf of Mexico ocean dynamics. Frontiers in Marine Science 28 (11), 1400560, doi: https://doi.org/10.3389/fmars.2024.1400560.

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.

National Academies of Sciences, Engineering & Medicine (2018). Understanding and predicting the Gulf of Mexico loop current: critical gaps and Recommendations, The National Academies Press, Washington, DC (2018). 10.17226/24823.

Oey L-Y, Ezer T, Foristall G, Cooper C, DiMarco S & Fan S (2005a). An exercise in forecasting loop current and eddy frontal positions in the Gulf of Mexico. Geophysical Research Letters 32, L12611. doi: https://doi.org/10.1029/2005GL023253.

Oey L-Y, Ezer T & Lee H-C (2005b). Loop Current rings and related circulation in the Gulf of Mexico: A review of numerical models and future challenges. In: Sturges, W.T. & Lugo-Fernandez A. (Eds.), Circulation in the Gulf of Mexico: observations and models, Geophysical Monographs series, 161, American Geophysical Union, Washington DC (2005), pp. 31-56.

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.

Sharma N, Storie JS, Obenour KM, Leber MJ & Srinivasan A (2016). Loop current hyperactivity: analysis of in-situ measurements in the Gulf of Mexico. Offshore Technology Conference, Houston, Texas, USA, 2-5 May 2016. doi: https://doi.org/10.4043/27229-MS.

Schmitz, WJ (2005). Cyclones and westward propagation in the shedding of anticyclonic rings from the loop current. In: Sturges, W.T. & Lugo-Fernandez, A. (Eds.) Circulation in the Gulf of Mexico: observations and models, Geophysical Monographs series, 161, American Geophysical Union, Washington DC (2005), pp. 31-56.

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.

Sonnewald 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.

Srinivasan A, Chin TM, Chassignet EP, Iskandarani M & Groves N (2022) A statistical interpolation code for ocean analysis and forecasting. Journal of Atmospheric & Oceanic Technology, 39(3), 367-386. doi: https://doi.org/10.1175/JTECH-D-21-0033.1.

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.

Steele ECC, Chen J, Ashton I, Pillai A, Jaramillo S, Leung P & Zarate L (2024) A spatiotemporal machine learning framework for the prediction of metocean conditions in the Gulf of Mexico. Offshore Technology Conference, Houston, TX, USA, 6-9 May 2024. doi: https://doi.org/10.4043/35104-MS.

Storie J, Ramos R, Leber M, Nowak H, Ivanov L, Magnell B & Young M (2023). Evaluation of loop current / loop current eddy fronts to guide offshore oil & gas operations. Offshore Technology Conference, Houston, Texas, USA, 1-4 May 2023. doi: https://doi.org/10.4043/32643-MS.

Sturges WT & Lugo-Fernandez A (Eds.) (2005). Circulation in the Gulf of Mexico: observations and models, Geophysical Monographs series, 161, American Geophysical Union, Washington DC, pp. 31-56.

Vaughan A, Markou S, Tebbutt W, Requeima J, Bruinsma WP, Andersson TR, Herzog M, Lane ND, Chantry M, Hosking JS & Turner RE (2024) Aardvark weather: end-to-end data-driven weather forecasting. arXiv, 2404.00411. doi: https://doi.org/10.48550/arXiv.2404.00411.

Appendix 1

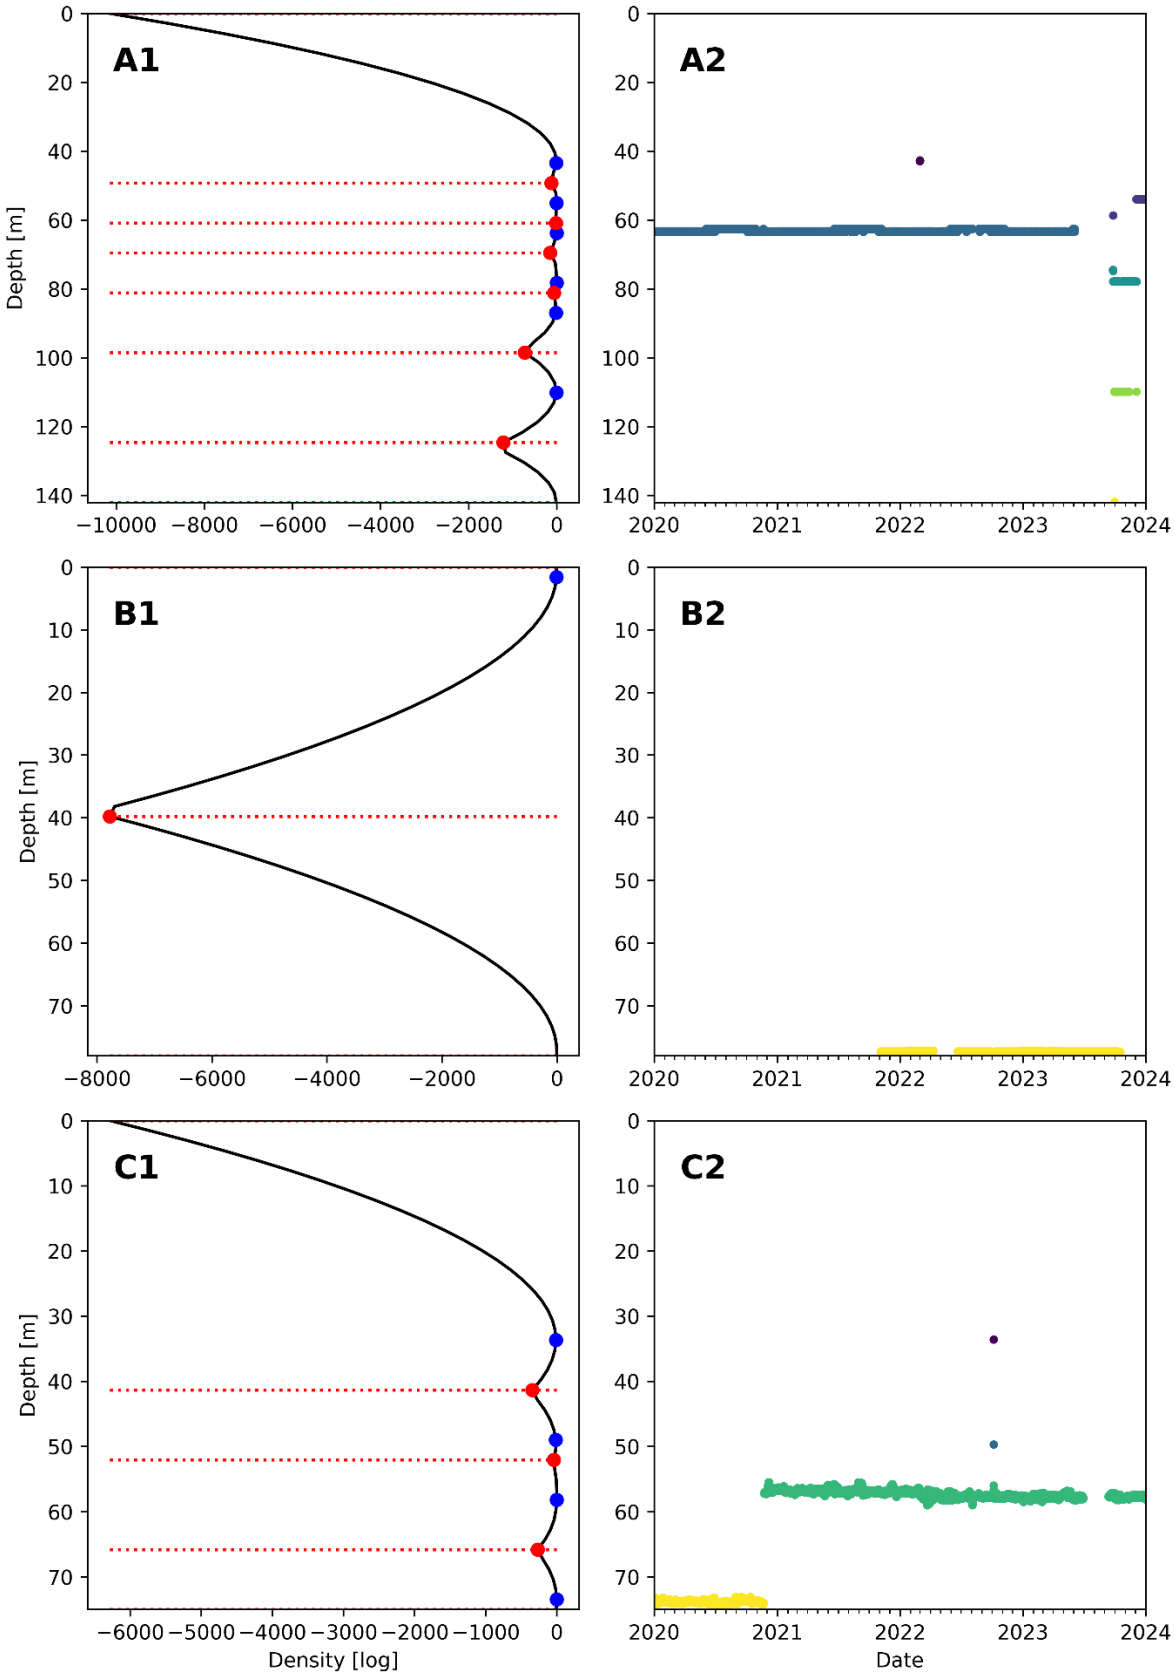

Figure A1: Demonstration of a one-dimensional Kernal Density Estimation (KDE)-based clustering method (left-hand column) used to ensure selection of the most appropriate data (longest continuous timeseries of constant depth) from the ADCP measurements (right-hand column) for each of the three input ADCP locations (top row: site 42390; middle row: site 42402; bottom row: site 42400).