Proceedings of the ASME 2023 42nd International Conference on Ocean, Offshore and Arctic Engineering
OMAE2023
June 11-16, 2023, Melbourne, Australia

OMAE2023-103020

ARE ROGUE WAVES PREDICTABLE FROM FIELD MEASUREMENTS?

Thomas Breunung
Department of Mechanical Engineering,
University of Maryland, College Park
College Park, United Sates

Balakumar Balachandran
Department of Mechanical Engineering,
University of Maryland, College Park
College Park, United States

ABSTRACT

Rogue waves, which are defined as waves with a wave height, or alternatively a crest height, exceeding the significant wave height by a certain factor, continue to endanger ships and offshore infrastructure. Hence, reliable rogue wave forecasting is of utmost importance to increase the safety for maritime operations. While the occurrence of rogue waves is widely acknowledged, their emergence remains unpredictable due to the lack of a well-accepted basis for explaining their occurrence. In fact, two popular mechanisms explaining the formation of rogue waves lead to considerably different conclusions about their predictability. On the one hand, a rogue wave could be formed by a superposition of wave trains with unknown phases. With this generation mechanism, rogue wave prediction is not viable. On the other hand, nonlinear focusing leading to the Benjamin-Feir instability gives rise to slowly developing rogue waves. Hence, this rogue wave formation could be detected with significant advance time. Given this background, there is an imperative need to address the basic question: Are rogue waves predictable?

In this article, the authors explore the predictability of rogue waves by constructing and parameterizing neural networks. The networks are trained on available buoy data, which allows not only for an assessment under the most realistic conditions but also for indicating the sufficiency of current ocean measurements for rogue wave prediction.

Keywords: Extreme waves, ocean buoys, machine learning, wave forecasting, extreme events

1. INTRODUCTION

Rouge waves are extreme waves, which are significantly higher than the surrounding waves. Such waves have damaged ships and offshore infrastructure and severely injured sailors as well as passengers [1,2]. To alleviate this danger, reliable rogue wave forecasts are of paramount importance.

Rogue waves can be generated by superposition of many elementary, linear or nonlinear, waves [2,3]. To observe such a constructive interference, the phases of the individual waves need to be aligned appropriately. However, from many wave models, including the state-of-the-art, operational spectral wave models [4,5] one cannot obtain this phase information and a random distribution of the phases is assumed. It has been argued that the initial conditions for these ocean models are inaccessible, and therefore, no meaningful phase information can be extracted [6]. Thus, the emergence of individual rogue waves remains unpredictable in this setting. Indeed, resolving the phase information has been identified as crucial for the prediction of rogue waves [7].

Irrespectively of the missing phase, quantities deduced from spectral wave models have been utilized for rogue wave predictions. For example, the excess kurtosis has been identified as a possible rogue wave indicator [8] and implemented in an operational wave forecasting system [9]. Alternatively, the Benjamin-Feir index deduced from the spectrum can indicate rogue waves resulting from the Benjamin-Feir instability [10]. Moreover, average sea statistics such as significant wave height, peak period, or skewness have been linked to rogue wave occurrences [11,12,13]. However, recent and extensive analyses [14,15] on available buoy data indicate that these parameters only poorly correlate with rogue wave occurrence in reality. In one study [15], the authors singled out the crest-trough correlation as the best predictor for certain rogue waves. However, a clear assessment of the predictable capabilities of the crest-trough correlation is not available in the literature.

Another popular mechanism for explaining the formation of rogue waves is the Benjamin-Feir instability [2,3,16,17]. It can be shown that the amplitude of a single wave train becomes unstable when appropriate modulations are added. Under certain conditions these side bands grow and a rogue wave is generated.

As this instability develops on a slow time scale, in principle, it could be detected in advance. Therefore, rogue waves generated by the Benjamin-Feir instability could be predicted with significant advance time. However, the Benjamin-Feir instability is commonly observed in highly idealized scenarios such as unidirectional, stationary and narrow band waves. Hence, many researchers have questioned the relevance of the Benjamin-Feir instability in realistic scenarios [15, 18, 19].

In summary, the two most common rogue mechanisms, superposition of elementary waves with random phases and the Benjamin-Feir instability differ considerably on the predictability of rogue waves. Moreover, with realistic data sets, the performance of the proposed rogue wave indicators is poor. This discrepancy illustrates the need to address the following fundamental question: Are rogue waves predictable?

In this article, the authors answer this basic question by assembling models, which can be used to predict upcoming rogue waves minutes in advance from field measurements. To this end, the authors construct and parameterize universal function approximators, more specifically, neural networks. The authors use available buoy data from the Coastal Data Information Program (CDIP), Scripps Institution of Oceanography [20]. By employing field measurements, a satisfactory answer is obtained under realistic conditions, and furthermore, it is examined whether current ocean wave observations are sufficient to predict rogue waves.

Fundamentally, rogue waves are predictable if there is a functional relationship between the waves that have passed prior to the rogue wave and the rogue wave event. If this functional relationship exists, it can be approximated by universal function approximators. Thus, universal function approximators can reveal the predictability of rogue waves.

In this work, the authors employ neural networks as universal function approximators to uncover a functional relationship for rogue wave predictions. With their universal function approximation capabilities [21], these networks are capable of capturing any functional relationship. Hence, these networks have the capability for predicting rogue waves, if they are predictable. Moreover, with this non-parametric model choice, the necessity to rely on underlying modelling assumptions is alleviated. The neural networks are fitted to available buoy data [20]. Thereby, the predictability is not only investigated under realistic conditions but it is also verified that current ocean wave observations are sufficient to predict rouge waves.

The general structure of the prepared data sets examined for the predictability of rogue waves is illustrated in Figure 1. Therein, two forty-minute sets of sea surface measurements are depicted. Only in the bottom time series, a rogue wave has occurred at about 35 minutes (highlighted in red), whereas no extreme wave occurred in the measurement data shown in the top of Figure 1. Now, the challenge is to predict upcoming rogue waves, based on measurements taken *prior* to the rogue wave event. More specifically, with a twenty-minute measurement

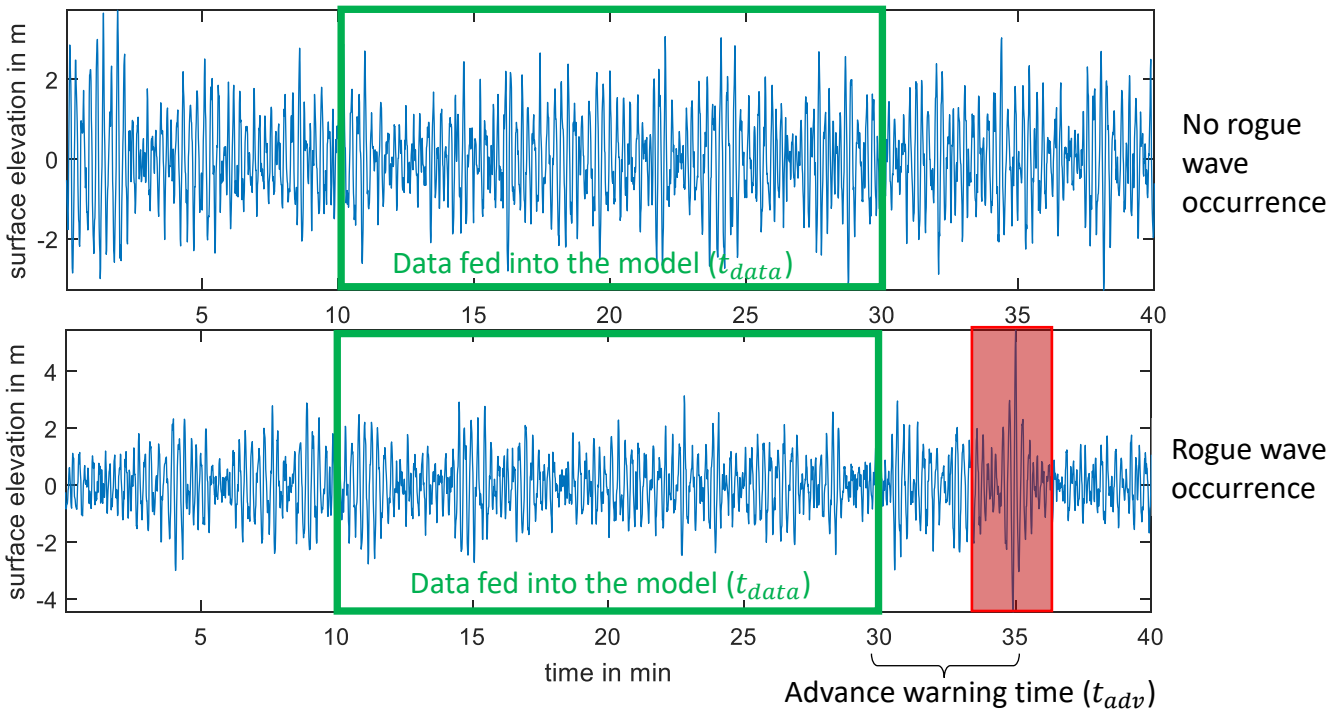

Figure 1: General structure of the prepared data sets. No rogue wave occurred in the measurement data shown in the top, whereas a rogue wave occurred at about 35 minutes shown in the measurement data shown in the bottom.

2. MATERIALS AND METHODS

available (cf. data inside the green box in Figure 1), the goal is

to determine whether a rogue wave will occur within a given advance warning time (five minutes in Figure 1) or not. Thus, the two twenty-minute measurement sets inside the green box in Figure 1 can be split into the two categories i) rogue wave occurs after the recording (bottom measurement) and ii) no rogue wave occurs after the twenty-minute data window (top measurement). Thus, rogue wave prediction can be formulated as a classification task, more specifically, a time series classification problem. Therein, the challenge is not to find appropriate categories, as the data is naturally labelled as 'rogue wave' and 'no rogue wave', rather the decision to which category a new measurement with unknown label belongs.

In the following, recurrent neural networks are utilized. To parameterize and validate these networks available buoy data [20] is utilized. This buoy data is described in the next section. Subsequently, the quality control of the data set, the data set preparation and the utilized neural network architecture are described.

2.1 Buoy Data

The buoy data set has been obtained from the Coastal Data Information Program (CDIP), Scripps Institution of Oceanography [20]. Overall, in this program, measurements from 166 buoys have been stored. As shown in Fig. 2, these buoys are primarily deployed at the shores of the continental US. Other buoys are located in the Pacific and some in lakes and sounds. In this study, the authors focus on rogue waves in the ocean, and hence, the recordings from inland buoys are not further considered. More specifically the CDIP buoys with the numbers 175, 177, 204, 205, 221, 230, 248, 251, and 253 are excluded from the presented analyses. With attention to the water depth variation, measurements from coastal as well as deep water

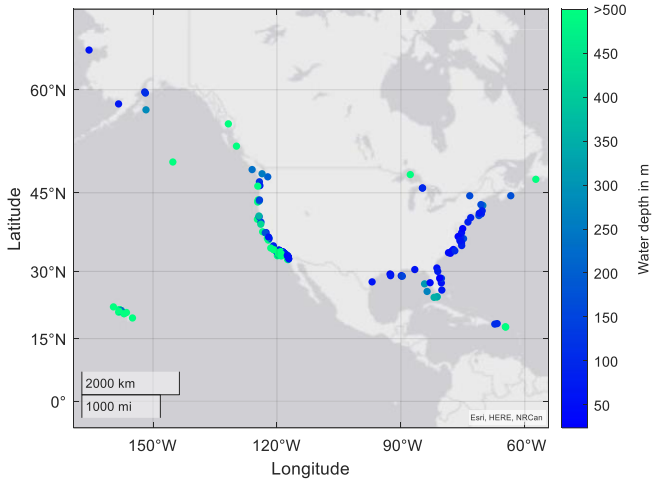

Figure 2: Location of the ocean buoys [20]. The color indicates the water depth. Additional buoys (not shown) are located near Floripa (Brazil), Anuu (American Samoa), Saipan (US), Guam (US), Palau (Marshall Islands) and Majuro (Marshall Islands).

are considered.

The CDIP-buoys are either Datawell directional waveriders MkIII [22] or Datawell directional waveriders MkIV [23] and in total the measurements from all buoys comprise of 660 GB of data. Amongst other quantities, these buoys have been used to record the buoy's vertical displacements with a sampling rate of 1.28 Hz for the MkIII system and 2.56 Hz for the MkIV system. From the buoy's vertical displacement measurements, rogue wave occurrences are apparent and hence they will be utilized in the following (cf. also Figure 1). All time series from all buoys together comprise more than 20 billion samples ($20 \cdot 10^9$), or equivalently 16 million half hour time series or 880 years of continuous data.

Due to the isolated buoy locations (cf. Fig. 2), no information about the spatial wave propagation is available, and hence, no predictions about the space-time evolution of ocean waves can be obtained. To realize these predictions, high resolution spatial measurements of the sea surface elevation need to be available.

Before proceeding with processing and identification of rogue waves, the enormous amount of data needs to be quality controlled. The utilized quality control protocol is detailed in the next section.

2.2 Quality control

Due to the harsh environments at sea, accurate measurements of ocean waves remain challenging [24]. Hence it is very likely that the vast data set [20] will be corrupted by measurement errors, sensor failures, and other malfunctions. To sort out unreliable data and only consider realistic wave profiles, the following judicious quality control is adapted.

First, the measurements are divided into thirty-minute intervals, which is a time span commonly used to compute sea state parameters in oceanography [12]. For each window, the following quality flags were evaluated:

1. Manufacturer quality flag: Datawell's waverider buoys are equipped with an algorithm, which flags unreliable measurements. If a single value within the thirty-minute window is flagged as 'questionable' or 'bad', the whole recording is discarded.
2. Operator quality flag: CDIP personnel also routinely check the recorded buoy data for irregularities. These inspections are based on computed spectra and other sea state parameters. Raw time series, however, are not checked, and hence, not flagged by the CDIP staff. Nevertheless, to still benefit from the decades of experience of these operators and ensure a rigorous quality control, this quality control is utilized for the time series as follows. If any sample of the recording under consideration contributes to a spectrum, which did not pass CDIP's quality control, then the whole measurement is discarded.
3. Spike detection: Even after passing the first two steps of the quality control, unrealistic spikes are observed in the recoded sea surface elevation (cf.

Figure 3). These spikes can be detected by monitoring the rate of change of the sea surface elevation. The authors employ the threshold value from [9]

$$S_y = \frac{4\pi\sigma}{T_z}\sqrt{2\ln(N_z)}, \qquad (1)$$

where σ denotes the standard deviation of the sea surface elevation, T_z is the mean zero upcrossing period, and N_z is the number of zero upcrossings. If the absolute value of the rate of change, calculated as the difference between two subsequent samples divided by the sample period, exceeds the threshold (1) once, then the whole 30-mintue measurement is discarded.

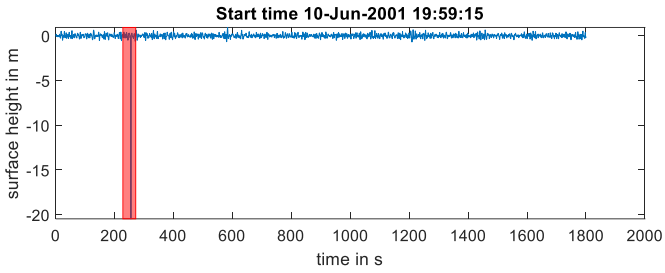

FIGURE 3: Unrealistic spike observed in the measured sea surface elevation.

4. Sensor range: Moreover, waves whose crest heights or trough depths exceeded the sensor range are found in the buoy data [20] (cf. Figure 4). Although, in principle, Figure 4 could show a rogue wave, the associated crest height exceeds the maximal sensor range of 20.47 m. This is considered to be indicative of a malfunction. Thus, measurements with sea surface elevations exceeding the sensor range are discarded.

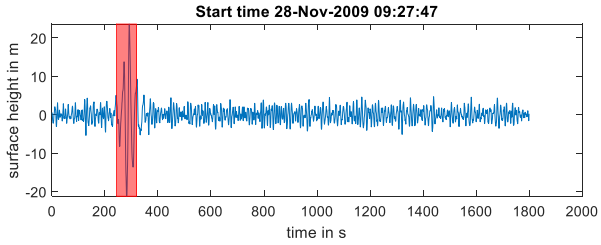

FIGURE 4: Sea surface elevation out of the sensor range.

5. Repetitive values: Additionally, in some measurements of the sea surface elevation, the same value occurs for multiple consecutive points (cf. Figure 5). Considering the ever changing, dynamic ocean surface, such behavior seems very unlikely. Thus, following the work in reference [12], every measurement with ten consecutive points of equal value is discarded.

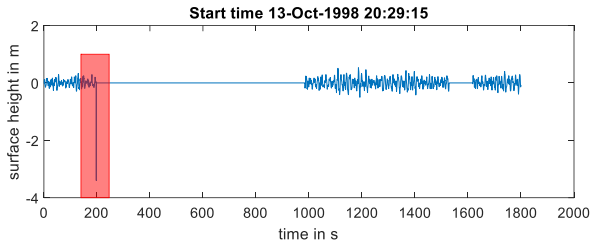

FIGURE 5: Consecutive points of equal value in the recoded sea surface elevation.

About ninety percent of the measurements of the original data set pass all five steps of the quality control. So, in total the buoy data [20] yields about 14 million quality controlled thirty-minute measurements. Subsequently, these recordings are scanned for rogue waves and organized into data sets, as described in the next section.

2.3 Data sets

If a thirty-minute measurement passes all previous five quality control steps, then, the measurement is scanned for rogue waves. From the literature (e.g., [3,11,12,14,15,18,25]), the following three rogue wave definitions are extracted

$$\frac{H}{H_s} > 2.2, \qquad \frac{H}{H_s} > 2, \qquad \frac{\eta_c}{H_s} > 1.25, \qquad (2)$$

where H denotes the wave height (trough to crest), η_c denotes the crest height, and H_s denotes the significant wave height. The significant wave height is defined as four times the standard deviation of the sea surface elevation (see, e.g., [3]).

If one of the criteria (2) is met, then the corresponding thirty-minute measurement of the sea surface elevation is included in a data set. Thereby, each criterion (2) is treated individually yielding a different data set. Moreover, the measurements are normalized so that the rogue wave event occurs at minute 25 within the saved recording. So, each saved recording consists of 25 minutes of measurement prior to the rogue wave and five minutes after the rogue wave. In addition to each sample with a rogue wave, another sample without a rogue wave from the same buoy location is saved. This procedure yields perfectly balanced data sets with an equal number of samples with and without rogue waves. Depending on the application, this balance could be altered in future studies.

As a final check, ten percent of the saved samples are selected for visual inspection and no systematic irregularities are detected. The authors note that the quality control described in Section 2.2 has been developed such that no systematic irregularities remain in the final data sets. Indeed, the quality control steps 3-5 (cf. Section 2.2) have been introduced after recognizing the irregularities shown in Figures 3-5.

A detailed overview of the three different data sets is provided in Table 1. Depending on the rogue wave definition (2), the final data sets consist of 40-400 thousand measurements of sea surface elevation. The rogue wave definition $H/H_s > 2.2$ yields the smallest data set, whereas about ten times more rogue waves are detected when the threshold is raised to two ($H/H_s > 2$). Owing to the definitions (2), the data set A is contained in the data set B. Furthermore, a large overlap is expected between the data set C and data set A as well as data set C and data set B. However, it is noted that data set C is not necessarily fully contained in data set B, due to the different rogue wave definitions (cf. equations (2)).

Table 1: Overview of the complied, quality-controlled data sets.

Data Set	Rogue wave Definition	Number of samples with rogue wave	Samples without rogue wave
A	$\frac{H}{H_s} > 2.2$	$18 \cdot 10^3$	$18 \cdot 10^3$
B	$\frac{H}{H_s} > 2$	$172 \cdot 10^3$	$172 \cdot 10^3$
C	$\frac{\eta_c}{H_S} > 1.25$	$27 \cdot 10^3$	$27 \cdot 10^3$

The data sets are prepared without storing the location or water depth of the ocean buoy. In subsequent studies, one could add these physical parameters and verify if the forecasting accuracy can be enhanced.

2.4 Neural Networks

In the past decade, neural networks have been successfully employed for various tasks such as image recognition, speech recognition, and genome sequencing. With their appealing promise to capture any underlying input-output relationship, assuming that enough data for parameter tuning is available, these networks are, at least in principle, capable of uncovering a hidden relationship between rouge wave events and waves prior to the rogue wave.

Within this work, recurrent neural networks with LSTM cells [26] are employed. Recurrent neural networks are designed to process time series and feature internal states. As a distinction, LSTM-networks can flexibly retain and erase their internal states. LSTM-networks have been successfully employed to handle complex time series [27] and ocean wave measurements [28,29].

A stacked architecture shown in Figure 6 is utilized, wherein each stage consists of the four layers. The first layer is a LSTM layer with N_{LSTM} hidden units. After that a batch normalization rescales each batch to have zero mean and unit variance. Subsequently, a dropout layer is used to set each feature to zero with a probability of p_D. Dropout layers are commonly employed to prevent overfitting, which often arises with neural networks [30]. Finally, a fully connected layer is used to extract N_f

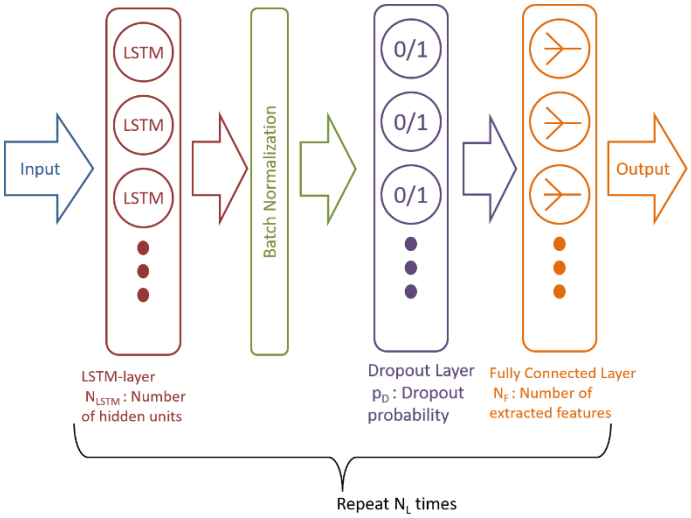

Figure 6: Employed architecture of the neural network.

features from the data. The four layers (LSTM, Batch normalization, fully connected and Dropout) are repeated N_L times.

The output of the network is a probabilistic classifier. For each class (i.e., "no rogue wave" and "rogue wave"), the output returned is a probability indicating the likelihood that the supplied sample belongs to the corresponding class. These two probabilities necessarily sum up to one. To yield a definite prediction, the class with the higher probability is selected as forecast.

The weights of the neural network are obtained via a stochastic gradient descent implemented within the Adam optimizer [31]. Therein, 80 percent of each data set (cf. Table 1) is utilized, with the remaining 20 percent being reserved for testing. The hyperparameters are selected upon best performance on the testing portion of the data set A. The final hyperparameter values are reported in Table 2. With the hyperparameter choice listed in Table 2, the network has about 4600 trainable parameters.

Table 2: Choice of hyperparameters

Parameter	Value
Hidden units N_{LSTM}	10
Dropout probability p_D	0.05
Number of features in dense layer N_f	2 (except 50 for the first one)
Number of stacked layers N_L	4

The obtained hyperparameters are verified with the state-of-the-art, automated hyperparameter tuning algorithm tune [32] and no significant improvements could be found. Moreover, alternative networks architectures, notably convolutional neural networks [33] and transformer networks [34] have also been tested. However, the depicted architecture in Figure 6 with the

hyperparameters listed in Table 2 yielded consistently the best results.

The networks are constructed and trained by using TensorFlow [35] (version 2.9.1). The training and testing experiments are conducted utilizing a NVIDIA Quadro P1000 GPU unit.

3. RESULTS AND DISCUSSION

The rogue wave predictions are obtained as follows. To train the neural network, twenty minutes of sea surface elevation are extracted from each sample ($t_{data} = 20$ min cf. Fig. 1). This window length yields time series with 1536 samples. For samples with a rogue wave the twenty-minute window is taken from the measurements *prior* to the rogue wave, whereby duration between the last included sample and the rogue wave event is denoted as advance warning time t_{adv} (cf. Figure 1). After extracting the measurements, they are randomly shuffled and 80% of the measurements are utilized to train the network, whereas the remaining 20% are reserved for testing the performance of constructed model on new data. The network is trained to distinguish measurements after which a rogue wave occurred from recordings where no rogue wave occurred afterwards. Thereby, the network can be used to obtain an advance warning of an occurring rogue wave, with a certain accuracy.

In the following, the data set A ($H/H_s > 2.2$) is employed within the forecasting experiments first. Subsequently, the performance of the neural network on the other two data sets (B and C) is investigated.

3.1 Data set A ($H/H_s > 2.2$)

First, the capabilities of the neural network to predict rogue waves are illustrated with the data set A. Within this data set, a rogue wave is defined as a wave with a wave height H exceeding the significant wave height H_s by a factor of at least 2.2. About 18 thousand rogue waves from the buoy data [20] satisfy this rogue wave definition. Compared to the alternative definitions (cf. equation (2)), the data set A consists of the least number of waves.

Setting the advance warning time of one minute yields the result shown in Figure 7. The depicted confusion matrix shows all outcome possibilities of the prediction experiment. In the first case, labelled as ``true negative'' the model's prediction is correct that no rogue wave will occur after the measurement. The percentage of 40.82% indicates that in about four out of five measurements, the prediction from the neural network is correct on no upcoming rogue wave. However, the field ``false positive'' with the percentage of 9.19% implies that about one out of five measurements is incorrectly labelled as measurement prior to a rogue wave. For the ``false negative'' field, the network prediction is no upcoming rogue wave, whereas in reality, a rogue wave does occur. The percentage of 10.68% implies that in about one out of five rogue waves is not predicted by the neural network. The ``true positive'' field, in turn, shows the number of accurately predicted rogue waves. Overall, the network prediction is correct for four out of five rogue waves.

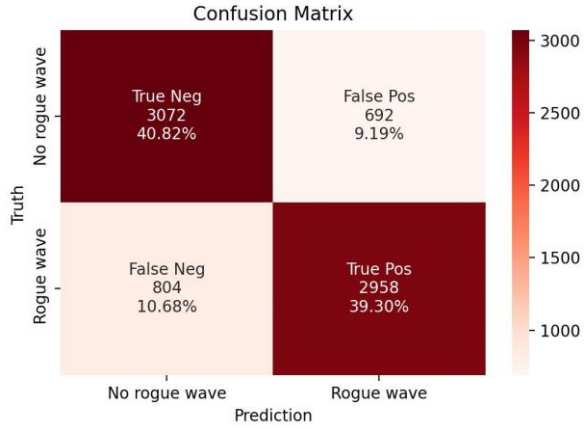

Figure 7: Prediction results for data set A ($H/H_s > 2.2$), with an advance warning time of one minute ($t_{adv} = 1$ min) and $t_{data} = 20$ minutes (cf. Figure 1).

In summary, from Figure 7, one can gather that rogue waves are indeed predicable from current ocean wave measurements. More precisely 80% of the rouge waves are predicable about one minute in advance. The overall accuracy calculated as the sum of the true positives and true negatives of the neural network's prediction is 80%. This implies that four out of five predictions of the neural network are correct. In total, about 3000 rouge waves are correctly predicted. To the best of the authors' knowledge, this is the most extensive rogue wave prediction experiment on field measurements that has been carried out.

The longer the advance warning times is, the more time is gained to enforce safety or undertake other countermeasures. Thus, it is natural to strive to extend the advance warning time. Subsequently, the advance warning time is increased from one minute to a maximum of five minutes and the forecasting experiments are repeated. The resulting percentages of true negatives, false positives, false negatives, and true positives are shown in Figure 8. Moreover, the percentage of correct predictions calculated as the sum of true positives and true negatives is included as accuracy. The overall trend is clear. As expected, the accuracy is found to deteriorate as the advance warning time is increased. Interestingly, the percentage of correctly predicted rogue waves remains approximately constant, whereas the correctly predicted no-rogue-wave samples decreases with increasing advance warning time. Even for the longest advance warning time of five minutes seven out of ten predictions are correct, or equivalently 2900 rogue waves

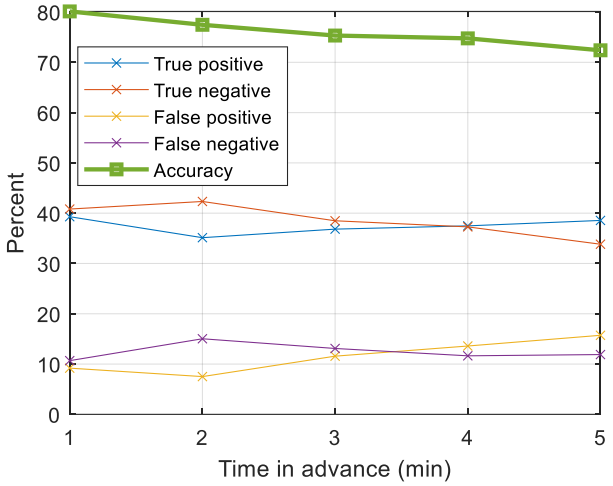

Figure 8: Prediction results for data set A ($H/H_s > 2.2$), for various advance warning times and $t_{data} = 20$ minutes (cf. Figure 1).

have been accurately predicted, while missing only 900 rogue waves.

The results shown in Figure 8 confirm the observation made in Figure 7. Rogue waves are predictable with a high accuracy of 70 to 80 percent within an advance warning time of several minutes. In the next section, it is investigated if that observation can be confirmed for the alternative rogue wave definitions utilized to compile the data sets B and C.

3.2 Data sets B ($H/H_s > 2$) and C ($\eta_c/H_s > 1.25$)

To verify the predictability of rouge waves satisfying alternative definitions, the forecasting experiments detailed in Section 3.1 are repeated for the data sets B and C. The data set B consists of about 10 times more rogue waves than data set A, while data set C is about twice as big as data set A.

Training and testing the LSTM-network on the data B for an advance warning time of one-minute yields the confusion matrix shown in Figure 9. For the advance warning time of one minute the percentage of correct predictions is about 76.4% and approximately 27 thousand rogue waves are correctly predicted from field data for the data set B.

The performance of the neural network on the data set B is slightly decreased compared to results shown with the data set A (cf. Figure 7). More precisely, the accuracy for data set B of 76.4% is less than the accuracy of 80.1% obtained for data set A. However, it is noted that this discrepancy could be explained with the hyperparameter choice (cf. Table 2). The hyperparameters were tuned for data set A but not data set B. Indeed, it is expected that for the more extensive and hence more complex data set B, a neural network featuring more parameters could help increase the accuracy further. Due to the excessive training time of more than a day with the current computational infrastructure, hyper parameter tuning for data set B was not carried out.

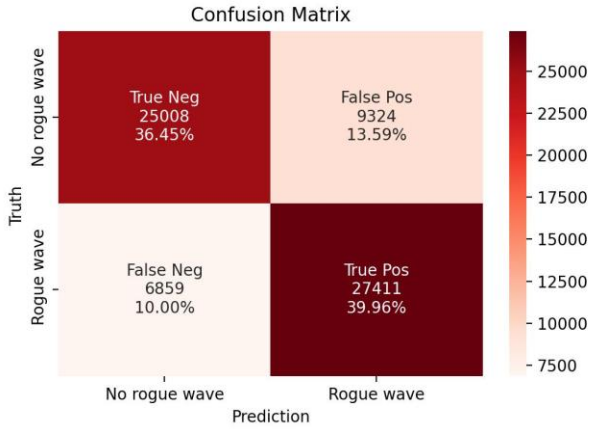

Figure 9: Prediction results for data set B ($H/H_s > 2$), with an advance warning time of one minute ($t_{adv} = 1$ min) and $t_{data} = 20$ minutes (cf. Figure 1).

In the same vein, the results for data set C are shown as a confusion matrix for an advance warning time of one minute in Figure 10. For this data set about 75.3% of all predictions are correct and 4251 rogue waves are correctly predicted in advance.

The percentage of correct predictions for data set C is found to decrease by about 5% compared to the data set A (cf. Figure 7). From this difference, one could conclude that rogue waves with large wave heights (i.e., $H/H_s > 2.2$) are easier to predict than rogue waves with large crests (i.e., $\eta_c/H_s > 1.25$). However, it remains to be explored how statistically relevant the observed discrepancy is. Additionally, the authors remark that observations based on such forecasting experiments can merely indicate correlations but not causations.

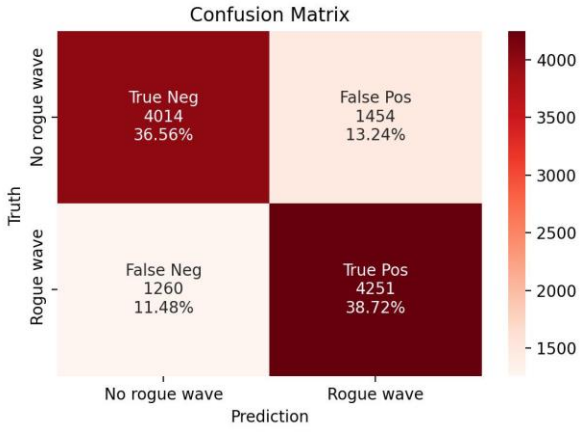

Figure 10: Prediction results for data set C ($\eta_c/H_s > 2$), with an advance warning time of one minute ($t_{adv} = 1$ min) and $t_{data} = 20$ minutes (cf. Figure 1).

4. CONCLUSION

In this work, the authors show that rogue waves are predictable from field measurements. More precisely, rogue wave can be predicted with an accuracy of about 80% one minute in advance from buoy data [20]. The forecasting accuracy drops to about 70 percent, when the advance warning time is extended to five minutes. These observations are verified by correctly predicting the emergence of thousands of rogue waves.

While within this article, it has been shown that rogue waves are predicable, the total accuracy of the predictions could be improved in future work. To this end, the ever-growing data set [20], alternative measurements, or different networks, and possibly an ensemble of multiple networks, could be utilized. Moreover, tools from explainable AI could be used to obtain physical insights from the trained, black-box models.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the National Science Foundation Grant No. CMMI1854532. They are also thankful to Samarpan Chakraborty, Samuel Dipasqua, Soheil Feizi, Kayo Ide, and Aya Abdelsalam Ismail of the University of Maryland, College Park, for ongoing discussions on this work. The data [10] were available from the Coastal Data Information Program (CDIP), Integrative Oceanography Division, operated by the Scripps Institution of Oceanography, under the sponsorship of the U.S. Army Corps of Engineers and the California Department of Parks and Recreation.

REFERENCES

[1] Didenkulova, I. I., Slunyaev, A. V., Pelinovsky, E. N., & Kharif, C. (2006). Freak waves in 2005. *Natural Hazards and Earth System Sciences*, *6*(6), 1007-1015.

[2] Kharif, C., & Pelinovsky, E. (2003). Physical mechanisms of the rogue wave phenomenon. *European Journal of Mechanics-B/Fluids*, *22*(6), 603-634.

[3] Dysthe, K., Krogstad, H. E., & Müller, P. (2008). Oceanic rogue waves. *Annu. Rev. Fluid Mech.*, *40*, 287-310.

[4] Tolman, H. L. (1991). A third-generation model for wind waves on slowly varying, unsteady, and inhomogeneous depths and currents. *Journal of Physical Oceanography*, *21*(6), 782-797.

[5] The WAMDI Group (1988). The WAM model—A third generation ocean wave prediction model. *Journal of Physical Oceanography*, *18*(12), 1775-1810.

[6] Komen, G. J., Cavaleri, L., Donelan, M., Hasselmann, K., Hasselmann, S., & Janssen, P. A. E. M. (1996). *Dynamics and modelling of ocean waves* (p. 554).

[7] Alam, M. R. (2014). Predictability horizon of oceanic rogue waves. *Geophysical Research Letters*, *41*(23), 8477-8485.

[8] Mori, N., & Janssen, P. A. (2006). On kurtosis and occurrence probability of freak waves. *Journal of Physical Oceanography*, *36*(7), 1471-1483.

[9] Janssen, P. A. E. M., & Bidlot, J. R. (2009). *On the extension of the freak wave warning system and its verification* Reading, UK: European Centre for Medium-Range Weather Forecasts.

[10] Janssen, P. A. (2003). Nonlinear four-wave interactions and freak waves. *Journal of Physical Oceanography*, *33*(4), 863-884.

[11] Christou, M., & Ewans, K. (2014). Field measurements of rogue water waves. *Journal of physical oceanography*, *44*(9), 2317-2335.

[12] Baschek, B., & Imai, J. (2011). Rogue wave observations off the US West Coast. *Oceanography*, *24*(2), 158-165.

[13] Adcock, T. A., & Taylor, P. H. (2014). The physics of anomalous ('rogue') ocean waves. *Reports on Progress in Physics*, *77*(10), 105901.

[14] Cattrell, A. D., Srokosz, M., Moat, B. I., & Marsh, R. (2018). Can rogue waves be predicted using characteristic wave parameters?. *Journal of Geophysical Research: Oceans*, *123*(8), 5624-5636.

[15] Häfner, D., Gemmrich, J., & Jochum, M. (2021). Real-world rogue wave probabilities. *Scientific Reports*, *11*(1), 1-11.

[16] Benjamin, T. B., & Feir, J. E. (1967). The disintegration of wave trains on deep water Part 1. Theory. *Journal of Fluid Mechanics*, *27*(3), 417-430.

[17] Yuen, H. C., & Lake, B. M. (1980). Instabilities of waves on deep water. *Annual Review of Fluid Mechanics*, *12*(1), 303-334.

[18] Fedele, F., Brennan, J., Ponce de León, S., Dudley, J., & Dias, F. (2016). Real world ocean rogue waves explained without the modulational instability. *Scientific reports*, *6*(1), 1-11.

[19] Gemmrich, J., & Cicon, L. (2022). Generation mechanism and prediction of an observed extreme rogue wave. *Scientific Reports*, *12*(1), 1-10.

[20] CDIP Buoy data https://doi.org/10.18437/C7WC72

[21] Nielsen, M. A. (2015). *Neural networks and deep learning* (Vol. 25). San Francisco, CA, USA: Determination press.

[22] Datawell (2021). Datawell Waverider Manial DWR-MkIII, https://datawell.nl/wp-content/uploads/2022/09/datawell_manual_dwr-mk3_dwr-g_wr-sg.pdf, Accessed: 12/06/2022

[22] Datawell (2021). Datawell Waverider Manial DWR4 https://datawell.nl/wp-content/uploads/2022/09/datawell_manual_dwr4.pdf, Accessed: 12/06/2022

[24] Barstow, S. F., Bidlot, J. R., Caires, S., Donelan, M. A., Drennan, W. M., Dupuis, H., ... & Wyatt, L. R. (2005). *Measuring and analysing the directional spectrum of ocean waves* (p. 465). Cost Office.

[25] Haver, S. (2001). Evidences of the existence of freak waves. In *Rogue waves* (pp. 129-140).

[26] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, *9*(8), 1735-1780.

[27] Vlachas, P. R., Byeon, W., Wan, Z. Y., Sapsis, T. P., & Koumoutsakos, P. (2018). Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *474*(2213), 20170844

[28] Breunung, T., & Balachandran, B. (2022, June). Freak Wave Forecasting: A Data-Driven Approach. In *International Conference on Offshore Mechanics and Arctic Engineering* (Vol. 85864, p. V002T02A002). American Society of Mechanical Engineers.

[29] Breunung, T., & Balachandran, B. (2022). Data-driven, high resolution ocean wave forecasting and extreme wave predictions. *Ocean Engineering,* 268, 113271.

[30] Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding machine learning: From theory to algorithms*. Cambridge university press.

[31] Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[32] Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J. E., & Stoica, I. (2018). Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*.

[33] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.

[34] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.

[35] Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... & Zheng, X. (2016). {TensorFlow}: a system for {Large-Scale} machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)* (pp. 265-283).