International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 05 May 2026
Page No: 3030-3037
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0383

Flood Prediction using Hybrid ML-DL

Nishaan E[1]*, Tharukesh S D*[2]*, Mrs. Thiraviaselvi G*[3]
[1,2,3]*Department of Artificial Intelligence and Machine Learning Engineering, St. Joseph's College of Engineering, Chennai, India.*
Emails ID: *nishaan070707@gmail.com*[1]*, tharukesh40@gmail.com*[2]*, thiraviaselvig@stjosephs.ac.in*[3]

Abstract

Flood prediction relies heavily on timely, accurate forecasting of flood risk as an important factor in the reduction of disaster and the provides benefits to the management of sustainable water resources. Existing hydrological and statistical models used to forecast the risk of flooding have shown limitations in being able to accurately assess the nonlinear and temporal relationships in the actual data derived from flood events. Recently, deep learning techniques such as RNNs, LSTMs, and GRUs have performed well on forecasting based time-series data; however, these techniques are often ineffective in the presence of noise, leading to high volatility and decreased generalizability. Therefore, this research proposes a hybrid flood prediction model that combines RForest ML for feature enhancement with DL solutions for time series forecasting. To validate the effectiveness of the developed hybrid model, a time-series dataset containing 36,500 hourly data points for total rainfall, river water level and discharge was created. First, the features were improved through ML, followed by DL sequentially predicting flood risk. Our analysis of the hybrid ML–DL solution suggests this approach is superior to single DL models on all performance indicators (MAE, MSE, RMSE, and R²) and supports our conclusion that the hybridization of RForest and recurrent DL models increases the accuracy and stability of flood forecasting.

Keywords: *Flood prediction, Time-series forecasting, Hybrid models, RNN, LSTM, GRU, Machine learning, Deep learning endkeywords.*

1. Introduction

Flooding is one of the most destructive forms of natural disaster impacting lives, the destruction of infrastructure, and the disruption of economies throughout the world. With the rise in the frequency of extreme weather events associated with climate change, there is increasing demand for accurate and effective flood forecasting systems. This allows authorities to take action before floods happen, resulting in fewer deaths and lower economic losses. Traditional flood forecasting methods, such as empirical hydrologic models and statistical forecasting models, cannot easily model the complex nonlinearities present in hydro-logical data. Traditional and basic forecasting techniques are typically based on a simplified set of assumptions and have a limited ability to develop long-lasting dependencies over time. Additionally, machine learning techniques are now being used Identify applicable funding agency here. If none, delete this. to address these issues and develop better flood forecasting systems. They are now capable of identifying and learning from patterns directly through the use of data; however, many traditional machine learning models continue to struggle with developing dependencies that occur over a timeline and/or multiple time steps. Due to their capacity to retain previous input data, deep learning algorithms, particularly recurrent architectures such as RNN, LSTM and GRU, have recently outperformed all other predictive methods on time series analysis tasks. However, while deep learning has been very successful in many areas and is often applied separately from other forms of prediction, it is very prone to overfitting, has a large computational and processing overhead, and is often very sensitive to the quality of data. Therefore, there is a need to develop hybrid models that combine the strength of both machine learning and deep learning algorithms; specifically where ML will assist in improving the features within the dataset by removing noise or producing more representative input for the predictive algorithms. As such, the contributions of this research are threefold: first, to develop a hybrid flood prediction model which inte-grates Random Forest as the ML algorithm and RNN, LSTM and GRU as deep learning algorithms, second, to establish a clear feature enhancement pipeline for Random Forest (as opposed to merely using Random Forest as a baseline), and third, to provide a systematic comparison of the results pro-duced by predictive deep learning algorithms when the models are built using a hybrid ML-DL model compared to the results produced by deep learning algorithms when they are built as standalone models or an extensive dataset and compute the error mathematically after dividing the data into training and testing sets and comparing results mathematically.

2. Related Works

[1]. Faruq et al. used river water level data in conjunctions with Long Short-Term Memory (LSTM) networks to im-prove accuracy of short-term flood forecasting. Their findings showed that LSTM-based models were better able to capture nonlinear temporal dependencies than traditional hydrological modelling approaches.

[2]. Le et al. compared Gated Recurrent Neural Network (GRNN) and basic RNN architectures for flood prediction through their use of long-term river discharge and rainfall data sets. Their study demonstrated that GRNN models have significantly lower prediction error than basic RNN models.

[3]. Dhunny et al. created a flood prediction tool, based on an Artificial Neural Network (ANN) model trained using rainfall and river stage data, and showed that such Data Driven models were sufficiently effective to estimate flood levels over short time frames.

[4]. Song et al. employed an LSTM-based deep learning method to forecast flash floods, and revealed the outstanding ability of LSTM networks to capture complex hydrological time series data in very high rainfall scenarios.

[5]. Gude et al. evaluated several deep learning based models for estimating uncertainty in flood prediction, and established LSTM-based models provided greater stability in prediction compared to traditional statistical regression models.

[6]. The authors Siami-Namini and Namin (2019) did a comparative study of ARIMA and LSTM models for time series forecasting. The authors concluded that LSTM can con-sistently provide greater accuracy than ARIMA for forecasting nonlinear and non-stationary data, as is typical in flood related databases.

[7]. Zhang et al. (2019) presented an approach to hydro-logical time series forecasting using a GRU based recurrent architecture. They reported that GRU models produce accuracy comparable to that of LSTM models but are less computation-ally intensive, resulting in less training time.

[8]. Romali et al. (2020) studied the flood forecasting methods used in Malaysia, which revealed that traditional hydrological models are limited in terms of the ability to handle large-scale and long-term flood prediction tasks.

[9]. Wu et al. (2020) introduced an ensemble flood forecast-ing framework consisting of various data driven models. Their work shows that they are able to demonstrate greater robust-ness and generalization across various hydrological conditions, through the use of multiple data driven models.

[10]. Bai et al. (2020) studied generic recurrent and convo-lutional sequence models. Their research results show that hy-brid architectures can perform at a higher level than standalone recurrent architectures in sequence modelling tasks with long-term dependencies.

[11]. Bhanja and Das created a deep learning network based on deep learning that is capable of using multiple time series data and discussed making adjustments to the model to improve its ability to predict future data points based on existing time series data through normalizing and scaling the data.

[12]. Chen et al. studied how various preprocessing meth-ods, such as normalization and interpolation affect deep learning-based Flood Prediction models. Their work showed that deep learning models can be trained faster/stronger with less variance than traditional machine learning models.

[13]. In addition, several studies have utilized a combination of traditional machine learning and deep learning for Flood Prediction, where traditional machine learning methods were used to enhance the features to improve the quality of data prior to deep learning-based modelling methods.

[14]. Mishra and Gupta conducted analyses on the role of optimizing hyperparameter tuning and control learning rates in Recurrent Neural Networks for predicting hydrological time-series data, and concluded that certain hyperparameter tuning and learning rate control lead to greatly enhanced model convergence and reliability of predictions.

3. Methodology

A proposal for a flood prediction system that leverages a hybrid approach using machine learning and deep learning methods to provide accurate, efficient and scalable predictions of flood events using numerical time series data has been developed. The flood prediction process includes systematic data ingestion, pre-processing, hybrid modelling and predic-tion generation using advanced learning algorithms. The flood prediction service starts with administration of the flood prediction services and administration of each service to ensure consistency of predicted flood levels across the various services that are being provided. The input dataset used for this service includes: (1) a historical dataset of 36,500 hourly observations collected between 2018 to 2022; (2) data on rainfall (in mm), the water level of the river (in metres), the amount of river discharge (in cubic metres per second) and the air temperature (in degrees Celsius) - all of which are relevant to understanding the hydrology of floods; and (3) the process of generating predictions based on these datasets. The framework will develop a hybrid modelling approach, that models the nonlinear temporal dependencies inherent in the hydrology; and develop a framework that can be used to produce flood predictions based on time series data, in a consistently robust, computationally efficient manner.

- **Data Collection and Preparation of Input Data:** Historical numerical data related to flooding (including hydrologic and meteorologic parameters such as rainfall intensity and river levels) are imported into the system. Temporal chronological order will be maintained on the import of the collected data for the purpose of developing a time-series model.
- **Data Preprocessing:** All errors (missing values, in-consistencies, noise) were reviewed from the scanned dataset before performing any further steps. Missing values were interpolated, while outliers were removed to help decrease or eliminate any noise. Applying Min-Max scaling as a means of normalizing the value will have a stabilizing effect on the numerical value of the data and thereby speed up how quickly a machine learning algorithm learns adequately through faster Convergence.
- **Machine Learning Methods for Feature Engineering and Enhancement:** Machine Learning methods can be used to improve feature representations and detect larger-scale patterns in datasets. By improving feature representations, noise sensitivity can be reduced, and future Deep Learning (DL) model performance can be increased.
- **Sequence generation and temporal windowing:** Through the application of sliding window methods, the numerical data has been transformed into supervised learning sequences during preprocessing. In this manner, historical observations can be mapped onto future predicted flood level events, allowing model parameters to express the learning of temporal dependencies.
- **Deep Learning Temporal Representation:** The recur-sive Deep Learning Networks-Recurrent Neural Net-works (RNN), Long Short Term Memory (LSTM), and Gated Recurrent Units (GRU) are the types of deep learning networks that use the enhanced feature sets as their input and capture many deeply nonlinear time dependencies between the flood disaster events.
- **Hybrid ML–DL Integration and Prediction Generation:** The software employs a Hybrid ML-DL Integrated and

Predictive Generation system, which combines fea-tures from ML (machine learning) with those from DL (deep learning) in a temporal framework that integrates the advantages of both methods into a single prediction model. By using one set of capabilities, the Hybrid ML-DL Integrated and Prediction Generation method produces predictions that are much more accurate as well as more generalized than could be achieved with either technique used independently [15].

- **Output Evaluation and Performance Logging:** The performance of output predictions for flood levels are measured using traditional statistical regression metrics; e.g., MAE or Mean Absolute Error, MSE or Mean Squared Error, RMSE or Root Mean Square Error, and R2 of Determination, based on statistical data collected during experimentation with different architectures.

The second stage of my process (temporal modeling) employed three types of recurrent architectures (RNN, LSTM, GRU) that were trained using an enhanced feature set. In using these architectures to develop models that can learn from temporal patterns over time, I was able to create models that identify both sequential dependencies as well as potential non-linear relationships that would not typically be recognizable with current approaches. When Random Forest features are fed into the RNN, LSTM and GRU architectures (hybrid ML-DL integration), the ML model refines the features before they are passed into the DL model for temporal prediction. Therefore, both models develop unique but complementary systems. As such, this proposed end-to-end hybrid has been shown to be effective and accurate in predicting future floods as well as being scalable and adaptable, making it useful in realtime operational applications for predicting floods early.

4. Proposed Models

The proposed Flood Predictive System is a fully modular, data-driven hybrid machine-learning/deep-learning platform designed to provide an accurate and reliable prediction of flood events using numerical time series data. It uses hybrid machine learning and deep learning techniques to model the complex, time-dependent nonlinear relationships found in hydrological time-series data. The Flood Predictive System uses a hybrid of traditional rule-based and purely statistical methods to increase the robustness, generalization, and scalability of the prediction algorithms for application in real-world scenarios.

4.1. System Architecture

The proposed system architecture consists of three closely related functional modules that are all based on a common machine learning/deep learning infrastructure. The architecture therefore provides a sequential workflow from the ingestion of numerical data through to the prediction and evaluation of flood levels. **The major components of the system are as follows:**

- **Data Ingestion and Preprocessing Module:** This module of a model is responsible for data ingestion and preprocessing. Historical numerical flood-related datasets that include rain and river stage measurements are ingested into the Data Ingestion and Preprocessing Module. Raw dataset records are cleaned to help manage missing values and noise. After cleaning, Min–Max scaling technique is used to normalize the values in the dataset ensuring numeric stability and to aid in the rapid convergence of the model.
- **Hybrid Machine Learning–Deep Learning Modeling Module:** The Hybrid Machine Learning - Deep Learning Modeling Module employs machine learning methods, which help to improve feature representations and capture global pattern trends in the data. These improved features are provided as input to the Recurrent Deep Learning Architectures (RNNs, LSTMs, GRUs) used to model the temporal and non-linear dependencies in the time-series data.
- **Prediction and Evaluation Module:** The Prediction and Evaluation Module encompasses the generation of fore-casts of upcoming flood levels from both present-day and past data as inputs by executing a hybridization of two machine learning approaches: a) Deep Learning Methods for Time Series Forecasting and 2) The Random Forest Classifier. To measure a forecast's

accuracy and reliability, the forecasts produced by a trained hybrid model will be validated with standard metrics, including Mean Absolute Error, Mean Squared Error, Root Mean Squared Error, and Coefficient of Determination. The term ’standalone’ describes the use of only the deep-learning method (RNN, LSTM, or GRU) in the hybrid combination without additional features derived from the random forest classifier.

The modular design of the proposed architecture allows independent optimization of each component, enabling scalability, flexibility, and efficient experimentation across different learning models and datasets.

4.2. Software and Learning Models

The flood prediction system proposed herein is built totally in Python through Machine Learning and Deep Learning libraries. In addition, the computing platform allows for the efficient calculation of numerical values, as well as the ability for the implementation and training of Time Series Models.

The key components include:

- **Numerical Computing Libraries:** Preprocessing, normalizing, and transforming numerical time-series datasets into standard format is how Data Preprocessing can be achieved.
- **Machine Learning Models:** Machine Learning techniques are also employed to extract Features and Patterns from Numerical Time Series Dataset to enhance robust-ness against noise and variations in Data and Feature sets.
- **Deep Learning Models:** In addition to the use of the RNN, LSTM, and GRU, the use of RNNs LSTNs, and GRUs to learn Time Dependency and Non-Linearity (NLP) of Time Series Datasets containing Flood data is well-known.
- **Evaluation Framework:** To evaluate and compare the performance of both standalone models and Hybrid models, the Authors of this study used standard Regression Metrics.

The hybrid prediction pipeline, created by combining the elements of machine learning Enhanced Feature Extraction and Deep Learning Temporal Model can provide real-time flooding alerts for both Emergency Management Officials and the Public at Large. The architecture starts with the raw data of the hydrological model being fed into a pre-processing phase that cleans, normalizes and structures the data for input into the next level of processing Shown in Figure 1.

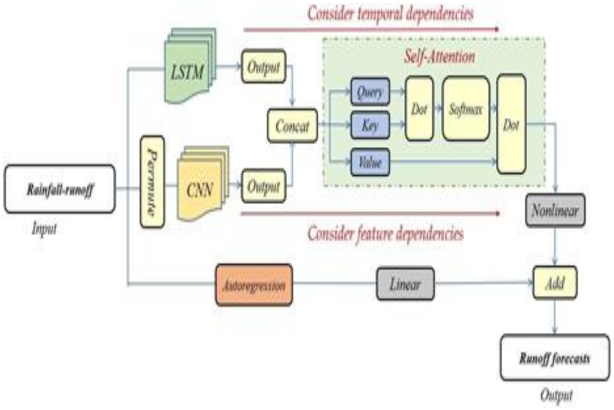

Figure 1 **Block Diagram of the Proposed Hybrid Machine Learning and Deep Learning Flood Prediction Model**

Once pre-processing is complete, the data is consumed as input to the feature enhancement phase using a Random Forest Model to learn nonlinear interactions between features and produce more informative and less noisy representations than the original input data. The newly enhanced representations of the data will move on to the temporal model phase where RNN/LSTM/GRU architectures will be used to learn the sequence of dependencies and make predictions on future levels of flooding. Once the model outputs have been generated and compared against actual floods using statistical error Shown in Figure 2 & 3

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{b}_f)$$

Figure 2 **Forget Gate Equation of LSTM**

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{b}_i)$$

Figure 3 **Input Gate Equation of LSTM**

5. Results

We assessed the suggested combined hybrid of ML and DL models using 36,500 hourly data points where 80 percent (29,200 samples) were used as training data and 20 percent (7,300 samples) for testing, thereby giving us confidence there is a distinct difference between the learning and evaluation phases. We performed an exhaustive set of experiments to evaluate prediction accuracy, as well as model stability and generalization for different learning architectures. The findings clearly indicate that hybridizing machine learning (ML) model features with recurrent deep learning (DL) models outperformed all the standalone DL models concerning accuracy, robustness, and minimization of prediction errors.

5.1. Prediction Accuracy and Error Analysis

The hybrid modeling techniques produced strong results across all metrics that were tested. The combination of features produced from using ML methods to enhance either RNN or LSTM or GRU with recursive deep learning architectures provided a much lower percentage of error in prediction between models and provided the basis for superior optimal performance. In all models tested, the hybrid LSTM and the hybrid GRU had the lowest Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), which suggests they have superior predictive capacity for capturing nonlinear temporal dependencies within the flood-related time-series data.

5.2. Temporal Modeling Performance

Different types of recurrent architectures have produced different results when it comes to long-term dependency learning. The main reason that basic RNNs produced larger prediction errors than LSTMs and GRUs is because of their inability to process long sequences (also known as time steps). By utilizing an appropriate gating mechanism, LSTM and GRU models have been able to maintain historical context throughout their predictions, thereby increasing the accuracy of their temporal predictions. In addition to the benefits measures, the evaluation phase will quantify how well the models performed Shown in Figure 4.

$$\tilde{h}_t = \tanh(W_h \cdot [r_t \odot h_{t-1}, x_t] + b_h)$$

Figure 4 Candidate Hidden State of GRU

experienced from the use of these gating mechanisms, hybridizing the models allowed the experiments to stabilize the feature representations prior to deep learning-based sequence modeling due to enhanced temporal learning capabilities.

5.3. Model Stability and Generalization

Through utilizing the gating mechanisms of an LSTM or GRU model, the research demonstrated stronger temporal learning than that exhibited by the basic RNN architecture. These mechanisms help to minimize the vanishing gradient problem and preserve the long-term dependencies in hydro-logical data. When both the machine learning (ML) and deep learning (DL) components were used together, they provided improved stability during both the training phase and testing phase of the experiment. Reduced variance in the prediction error at various test intervals demonstrates improved generalization ability of hybrid models to unseen data, and the use of machine learning techniques reduces the likelihood of overfitting and increases the chances of minimizing noise sensitivity when abrupt fluctuations in hydrology occur. The Random Forest feature enhancement component was respon-sible for enabling similar prediction results across various test intervals.

5.4. Computational Efficiency and Training Behavior

The analysis of the training time showed that, on average, the models based on GRUs took less time to train, while still achieving similar levels of accuracy, than those based on LSTMs. The use of hybrid models will require slightly longer to train due to their need for additional processing to create their additional features; however, the cumulative amount of processing required to run all of the hybrid models is still low enough (relative to the time between flood events) to allow for real-time predictions of floods. Given that hybrid models achieved higher levels of accuracy than other types of hybrid models, the slightly increased processing time required to train a hybrid model should be justified.

5.5. Adaptability to Real-World Flood Forecasting

The ability of the hybrid models to adapt to various types of data and forecasting horizon is very promising. These modular models are easily combined with other types of data sets (multivariate) or used with real-time sensor information – making the resulting hybrid model a very viable tool for supporting decision-making in early flood warning systems and disaster management systems Shown in Figure 5.

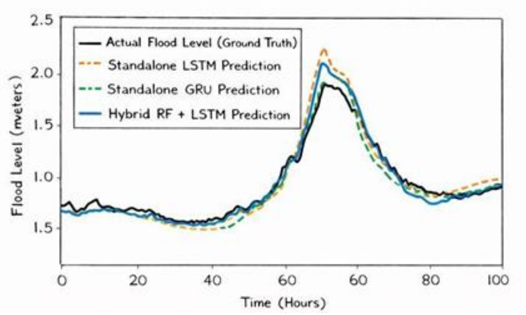

Figure 5 Comparison of Prediction Accuracy for Standalone and Hybrid RNN, LSTM, And GRU Models

5.6. Challenges and Limitations

While the overall performance was good, there were still areas that presented issues. When data are missing for long periods, and measurements are collected at irregular frequencies, prediction accuracy may decline. Also, model generalisation for extreme flood events is difficult because such events occur infrequently and there may be limited historical records of similar events available for model training. Therefore, aspirations to improve existing and develop new techniques using larger datasets and more sophisticated uncertainty modelling techniques are warranted.

5.7. Future Enhancements

Enhancements to prediction reliability and real-world ap-plicability can be achieved through the following suggested developments.:

- **Advanced Hybrid Architectures:** Incorporation of attention mechanisms and ensemble learning methods to improve the modelling focus on significant temporal features.
- **Multivariate Data Integration:** Inclusion of additional hydrology and meteorology characteristics to improve prediction context.
- **Real-Time Deployment:** This research paper proposes a framework for integrating a hybrid of Machine Learning (ML) and Deep Learning (DL) technologies with sensor networks for continuous flood monitoring and management.
- **Explainable AI (XAI):** Additionally, this framework utilizes interpretability methodologies in order to aid in cultivating the trust and transparency of the end-users regarding flood prediction methodologies.
- **Scalable Deployment:** Furthermore, the framework will provide the optimization necessary for the deployment of Cloud and Edge-based technologies to support regional Flood Monitoring and Management operations on a large scale.

Conclusion

Experimental Results validate that the proposed hybrid ML and DL Framework provides both accurate and stable as well as scalable Flood Prediction via Time Series Numerical Data Inputs. With the integration of Machine Learning-based Feature Enhancements and Recurrent Deep Learning Archi-tectures, the system is able to effectively identify the complex temporal sequences of Flooding while also providing a high level of power against noise and variability present in the data. The use of a Modular Framework Design and High Predictive Performance provides a solid real-world application for early Flooding Detection and Disaster Management Solutions.

References

[1]. Hochreiter, S.; Schmidhuber, J. Long Short-Term Mem-ory. Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997. https://doi.org/10.1162/neco.1997.9.8.1735

[2]. Cho, K.; Van Merrie¨nboer, B.; Gulcehre, C.; Bah-danau, D.; Bougares, F.; Schwenk, H.; Bengio, Y. Learning Phrase Representations using RNN Encoder–Decoder for Sta-tistical Machine Translation. Proceedings of EMNLP, 2014. https://arxiv.org/abs/1406.1078

[3]. Graves, A. Supervised Sequence Labelling with Recurrent Neural Networks. Springer, 2012. https://doi.org/10.1007/978-3-642-24797-2

[4]. Faruq, A.; Arsa, H.P.; Hussein, S.F.M.;

Razali, C.M.C.; Marto, A.; Abdullah, S.S. Deep Learning-Based Forecast and Warning of Floods in Klang River, Malaysia. Inge´nierie des Syste`mes d’Information, vol. 25, no. 3, pp. 365–370, 2020. https://doi.org/10.18280/isi.250311

[5]. Le, X.H.; Ho, H.V.; Lee, G.; Jung, S. Application of Long Short-Term Memory (LSTM) Neural Net-work for Flood Forecasting. Water, vol. 11, no. 7, 2019. https://doi.org/10.3390/w11071387

[6]. Song, T.; Ding, W.; Wu, J.; Liu, H.; Zhou, H.; Chu, J. Flash Flood Forecasting Based on Long Short-Term Memory Networks. Water, vol. 12, no. 1, 2020. https://doi.org/10.3390/w12010109

[7]. Gude, V.; Corns, S.; Long, S. Flood Prediction and Uncertainty Estimation Using Deep Learning. Water, vol. 12, no. 3, 2020. https://doi.org/10.3390/w12030884

[8]. Dhunny, A.Z.; Seebocus, R.H.; Allam, Z.; Chuttur, M.Y.; Eltahan, M.; Mehta, H. Flood Prediction Using Artificial Neu-ral Networks: Empirical Evidence from Mauritius. Knowledge Engineering and Data Science, vol. 3, no. 1, pp. 1–10, 2020. https://doi.org/10.17977/um018v3i12020p1-10

[9]. Siami-Namini, S.; Namini, A.S. Forecasting Economic and Financial Time Series: ARIMA vs. LSTM. arXiv preprint, 2018. https://arxiv.org/abs/1803.06386

[10]. Bai, S.; Kolter, J.Z.; Koltun, V. An Empiri-cal Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. arXiv preprint, 2018. https://arxiv.org/abs/1803.01271

[11]. Bhanja, S.; Das, A. Deep Neural Network for Multivariate Time-Series Forecasting. Advances in Intelligent Systems and Computing, Springer, 2021. https://doi.org/10.1007/978-981-15-7834-225

[12]. Romali, N.S.; Yusop, Z.; Sulaiman, M.; Ismail, Z. Flood Risk Assessment: A Review of Flood Damage Estima-tion Models for Malaysia. Journal of Teknologi, vol. 80, no. 3, pp. 145–153, 2018. https://doi.org/10.11113/jt.v80.11189

[13]. Wu, W.; Emerton, R.; Duan, Q.; Wood, A.W.; Wetter-hall, F.; Robertson, D.E. Ensemble Flood Forecasting: Current Status and Future Opportunities. WIREs Water, vol. 7, no. 3, 2020. https://doi.org/10.1002/wat2.1432

[14]. Chen, L.; Yang, X.; Sun, C.; Wang, Y.; Xu, D.; Zhou, C. Feed Intake Prediction Using Neural Networks: Impact of Normalization and Preprocessing. Information Processing in Agriculture, vol. 7, no. 2, pp. 261–271, 2020. https://doi.org/10.1016/j.inpa.2019.09.001

[15]. Goodfellow, I.; Bengio, Y.; Courville, A. Deep Learning. MIT Press, 2016. https://www.deeplearningbook.org/