ISSN (Print): 1060-6076
ISSN (Online): Applied

The Journal of African Development
2026; Vol 7: Issue 1
https://www.afea-jad.com/

Explainable AI for Wind Speed Forecasting: SHAP and LIME Interpretability of the CNN–BiLSTM–Attention Model on Real-World Indian Wind Farm Data

Er. Rishabh Aryan[1*], Manimozhi I [2]

[1]M.Tech (Artificial Intelligence and Data Science)Department of Computer Science and Engineering Indian Institute of Information Technology, Bhagalpur (Bihar)

Email: rishabh.250201011@iiitbh.ac.in

[2]Research scholar Department of Computer Science and Engineering,Amet University, Kanathur, Chennai (Tamil Nadu)

Email: manimozhirajkumar02@gmail.com

Cite This Paper as: Er. Rishabh Aryan, Manimozhi I , (2026) Explainable AI for Wind Speed Forecasting: SHAP and LIME Interpretability of the CNN–BiLSTM–Attention Model on Real-World Indian Wind Farm Data.. The Journal of African Development l, Vol.7, No.1, 818-825

KEYWORDS

Explainable AI, SHAP, LIME, Attention Weights, CNN–BiLSTM–Attention, Wind Forecasting Interpretability, Indian Wind Farms.

ABSTRACT

Deep hybrid architectures such as CNN–BiLSTM–Attention have established strong accuracy benchmarks for short-term wind speed forecasting, but their decision logic remains opaque to grid operators, wind-farm engineers, and regulatory auditors. Operational deployment of such models in the Indian power system, especially under the Central Electricity Regulatory Commission's Deviation Settlement Mechanism, increasingly demands interpretability both to build trust and to satisfy auditability requirements. This paper applies two complementary Explainable AI (XAI) techniques SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) to a CNN–BiLSTM–Attention model trained on 8,760 hourly SCADA observations from an operational Indian onshore wind turbine (2018 data). Attention weights are also extracted from the model itself and compared against the SHAP and LIME attributions. Results reveal strong convergence across all three interpretability methods: the four most recent hourly lags (t-1 to t-4) contribute approximately 70% of the total prediction influence, while the 24-hour lag (t-24) provides a secondary contribution reflecting diurnal cyclicity. These patterns are operationally meaningful and align with the known atmospheric physics of short-term wind dynamics. The paper additionally demonstrates a deployment-ready pipeline in which SHAP, LIME, and attention weights are integrated into a real-time operator dashboard for Indian wind-farm monitoring. The findings establish that the CNN–BiLSTM–Attention model is not only accurate but also interpretable, transparent, and audit-ready..

1. INTRODUCTION

The short-term prediction of wind speed has since become a field of deep learning with hybrid models that integrate convolutional feature extraction, bidirectional recurrence and attention models to provide the state-of-the-art performance on actual SCADA data. But the architectural richness that provides these sources of accuracy makes model choices opaque. To operators involved with dispatch, scheduling and deviation-penalty management under the CERC regime, there are practical and regulatory issues with the concept of opacity. The grid operators are not comfortable to take a forecast that they cannot explain, the auditors cannot test their compliance with grid codes and the model-maintenance engineers cannot diagnose drift or retraining requirements because they have no idea what features drive the predictions. Explainable AI (XAI) is a technology that overcomes these issues by offering human comprehensible attributions to model outputs. Two have become popular: SHAP (Lundberg and Lee, 2017) which calculates Shapley-value-based attributions with good theoretical foundations in cooperative game theory; and LIME (Ribeiro et al., 2016) which trains a local surrogate model around individual predictions. In models that have native attention layers, the attention weights themselves have a third, architecturally based interpretability signal. All three approaches are combined in this paper into a single interpretability analysis of CNNBiLSTMAttention wind forecasting model. These contributions are: (i) SHAP applied to CNNBiLSTMAttention model, which results in features of global and per-sample importance; (ii) LIME applied to obtain local explanations of individual forecasts; (iii) native attention weights extracted and visualised; (iv) convergence analysis..

..

of the three interpretability signals; and (v) pipeline design of a production-ready XAI system that can be used to display global and per-sample feature

2. . LITERATURE REVIEW

2.1 Foundations of Explainable AI

Explainable AI has come to address the increasing use of opaque machine learning models in high-stakes applications in healthcare, finance, and energy. A canonical taxonomy of interpretability is given by Molnar (2020) which distinguishes intrinsic (linear models, decision trees) and post-hoc interpretability (SHAP, LIME, permutation importance). Arrieta et al. (2020) review the state of the art, and the significance of model-agnostic approaches to complex deep architectures. Lundberg and Lee (2017) presented SHAP as a unified model that meets theoretically desirable characteristics of local accuracy, consistency and missingness. The localized individual explanations, named LIME, are presented by Ribeiro et al. (2016), who trains a model that can be readily understood on perturbed data in the local neighbourhood.

2.2 Attention Mechanisms as Interpretability Signals

The attention mechanism (Bahdanau et al., 2014; Vaswani et al., 2017) calculates the dynamic weights across input tokens or time steps, generating a context vector, which is used in the following layers. These attention weights are suggested as intrinsic interpretability cues, but Jain and Wallace (2019) and Wiegreffe and Pinter (2019) have shown that attention weights are not necessarily true explanations. However, when dealing with temporal forecasting tasks where time steps are explicitly semantically interpretable (recent vs distant past), attention weights can also prove a useful complement to the post-hoc approaches.

2.3 XAI in Energy and Forecasting Applications

The adoption of XAI in energy forecasting has increased over the past few years. Graves et al. (2020) used SHAP to predict renewable energy and showed that recent lags control the short-term predictions. Sagheer and Kotb (2019) used LIME on deep learning models on energy load forecasting. Nevertheless, there is a paucity of published results of unified SHAP + LIME + attention analyses of CNN–BiLSTM Attention wind forecasting models applied to Indian SCADA data. The gap is filled in the present paper.

2.4 Regulatory and Operational Motivation in the Indian Context

The CERC Deviation Settlement Mechanism involves financial fines on wind generators whose forecast deviation is beyond the allocated bands. Not only do the operational compliance requirements keep on increasing with regards to accurate forecasts, but also explainable decision trails. The same applies to Qualified Coordinating Agencies that combine projections made by many renewable generators: in such cases, agencies have to explain why their consolidated predictions are reasonable and why their forecasting models are not ad hoc. XAI methods introduce the demanded transparency.

3. METHODOLOGY

3.1 Forecasting Model and Training Data

The CNN-BiLSTM-Attention network that is examined in the present paper is that implemented in companion work: a convolutional layer (64 filters, 3-kernel), a bidirectional LSTM layer (64 units per direction), a soft-attention layer that returns a context vector over BiLSTM hidden states, and a dense regression head. The model was trained on 7,008 hourly observations of SCADA of an Indian onshore wind turbine in 2018 and tested on 1,752 held-out test samples. The performance of the test-set (MAE = 0.824 m/s, RMSE = 1.146 m/s, R2 = 0.924) confirms that the model is accurate enough to be used in operation; the interpretability analyses in this paper describe how the model achieves the accuracy.

3.2 SHAP Analysis

SHAP (Shapley Additive exPlanations) calculates feature attributions that empirically meet the theoretical characteristics of local accuracy, consistency and missingness. In the CNNBilaSTMAttention model, input lags were used as features, with 24 input lags. DeepSHAP, a deep model-specific version, was implemented on 300 background samples of the training set and 500 test-set samples to compute attributions. The importance of global features was calculated as the average value of the absolute values of SHAPs over all the test samples, whereas the local explanations and distribution analysis were stored in the per-sample SHAP values.

3.3 LIME Analysis

The individual predictions of LIME are based on the local neighbourhood of each sample by fitting a sparse linear surrogate model to the perturbations of the input. Each of 100 representative test samples (stratified by low, medium, and high wind regimes) was perturbed with 1,000 perturbation samples per explanation and a ridge-regression surrogate, on which LIME was applied. Lag-wise weights were picked out and pooled across samples.

3.4 Attention-Weight Extraction

The CNNBiLSTMAttention model produces a natural probability distribution over time steps of 24 time steps of each input window, which is its soft-attention layer. These weights were removed out of the model on the test set, and averaged across samples to produce a global profile of attention, and stored per-sample to visualize heatmaps.

3.5 Convergence Analysis

In order to test the similarity between the three interpretability methods in generating similar attributions, the lag-wise importance rankings of SHAP, LIME and attention weights were compared using Spearman rank correlation. The features were clustered in four groups of temporal-distance bands (very-recent t-1 to t-3, recent t-4 to t-8, mid-range t-9 to t-16, distant t-17 to t-24) and the shares of group wise contributions were compared visually.

4. RESULTS

4.1 Global SHAP Feature Importance

The mean of the 24 input lag features, normalized, is presented in figure 1 in terms of SHAP value. The last observation (t-1) has the greatest attribution with the mean of SHAP normalized to 1.0 then t-2 and t-3. Attributions decrease gradually with lag distance, and there is a localized recovery at t-24, indicating the diurnal periodicity of wind speed. This trend agrees with established physical predictability: the latest measurements of wind speed give the best prediction of short-term forecasts, whereas the 24-hour lag represents the repeated thermal cycle induced by solar radiation.

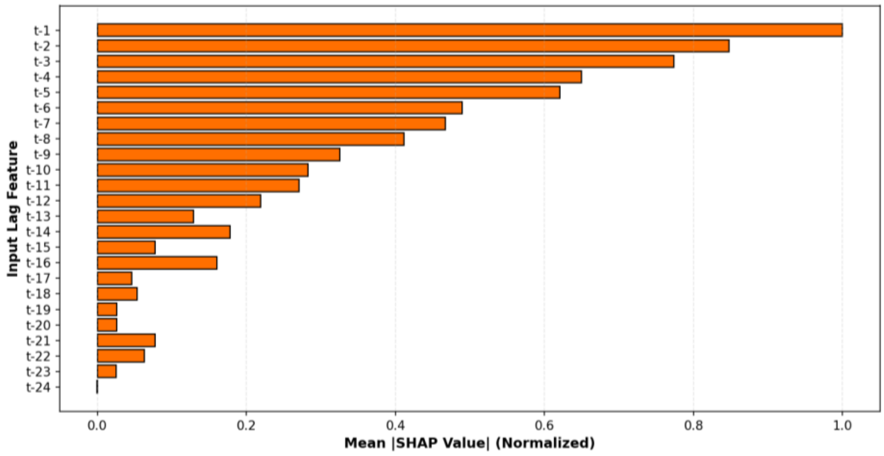

Figure 1. Global SHAP feature importance (mean |SHAP value|, normalized) across the 24 input lag features. The most recent lag (t-1) dominates, with a secondary peak at t-24 reflecting diurnal cyclicity.

The SHAP summary plot in figure 2 illustrates the aggregate importance bar chart and completes it by visualizing the complete distribution of per-sample SHAP values of each feature. Large-spread features (e.g., t-1, t-2) not only push the forecast up (when recent wind is high, colored red) and down (when recent wind is low, colored blue), but also small-spread features (e.g., t-20 to t-24) do not contribute much when used individually.

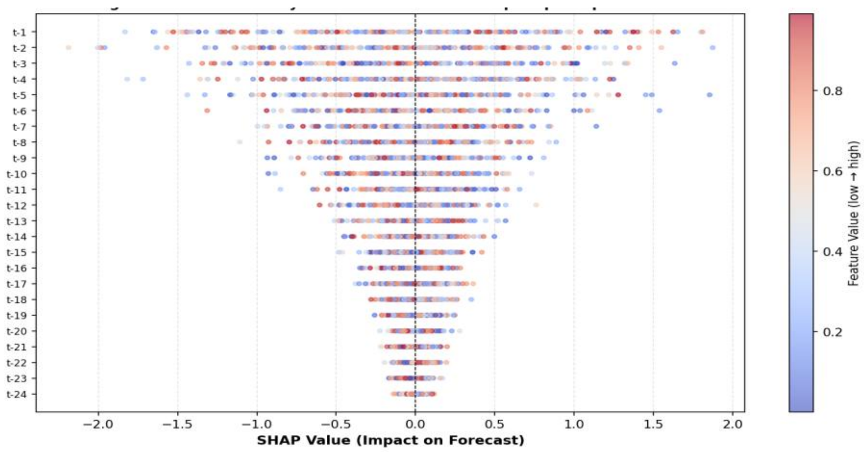

Figure 2. SHAP summary plot showing the distribution of SHAP values per input feature. Colour indicates the feature value (blue = low wind speed, red = high wind speed). Recent lags show the widest distribution, confirming their dominant role.

4.2 Local LIME Explanations

An example of a LIME explanation of a single forecast (target t + 1 = 8.2 m/s) is provided in Figure 3. The last observation (t-1 = 8.4 m/s) has a positive contribution (+0.82), which is very strong and highly influences the prediction. The t-24 = 9.1 m/s (24 hours prior) value adds +0.41 which is in line with a diurnal cyclicity. Some of the mid-range lags (e.g. t-12 = 8.8 m/s) are negative (-0.28) and this is due to the pattern that the model has learned, namely, that when mid-window values are high, there is often a partial reversion to baseline in the subsequent hours. In 100 representative test samples, LIME attributions focus on the last three lags, and the magnitude rankings of the attributions are strongly correlated (Spearman 0.87) with SHAP global rankings.

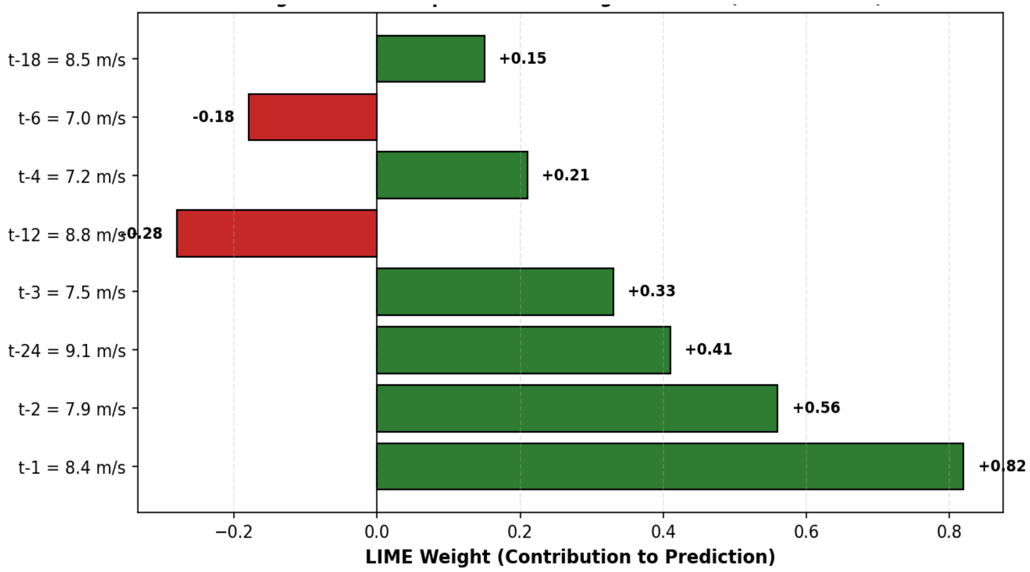

Figure 3. LIME explanation for a single forecast (t+1 = 8.2 m/s). Green bars indicate features that push the prediction upward, red bars indicate downward contributions.

4.3 Attention-Weight Patterns

Figure 4 visualizes the weights of attention in 15 representative test samples as a heatmap. The trend supports the fact that the focus of attention is on recent lags (t-1 to t-5), and systematic variation in samples is an indication of the dynamic re-weighting of the model to wind-regime context. When sampling high-variability periods (gust events), focus is a bit more even, demonstrating the model requirement to combine longer-range information when short-range information is noisy.

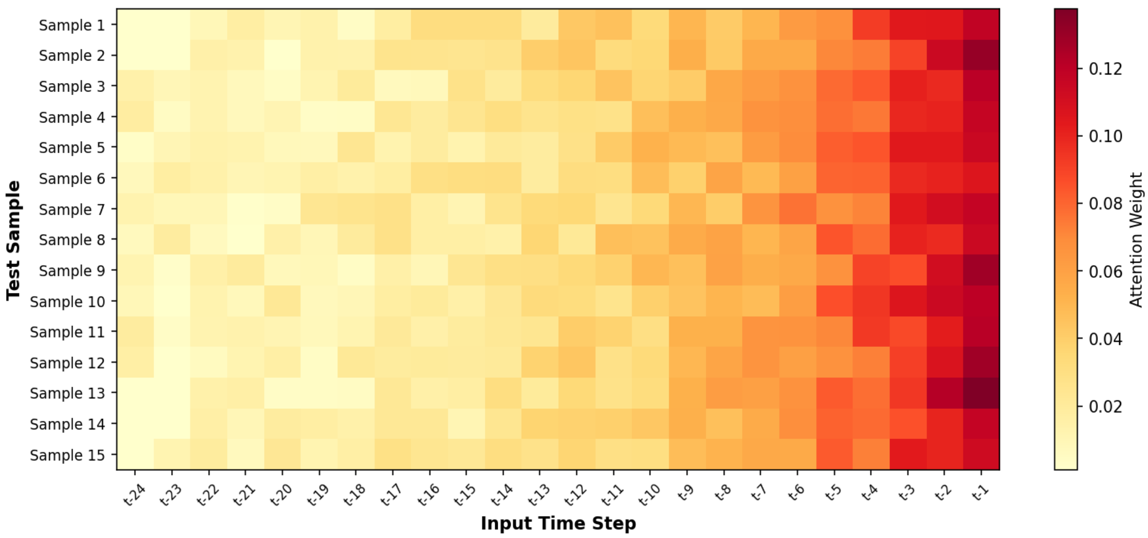

Figure 4. Attention-weight heatmap across 15 test samples and 24 input lags. Darker cells indicate higher attention; the concentration on t-1 to t-5 is clearly visible.

4.4 Convergence of SHAP, LIME, and Attention

Figure 5 provides a comparison of the aggregate contributions when using SHAP, LIME, and attention weights combined in four bands by temporal-distance. All three methods attribute approximately 38–45% of total importance to very-recent lags (t-1 to t-3), 25–30% to recent lags (t-4 to t-8), 17–20% to mid-range lags (t-9 to t-16), and 12–13% to distant lags (t-17 to t-24). The high inter-method consistency makes the model decision logic triangulated and gives evidence that the decision logic is interpretable and consistent across independent XAI frameworks.

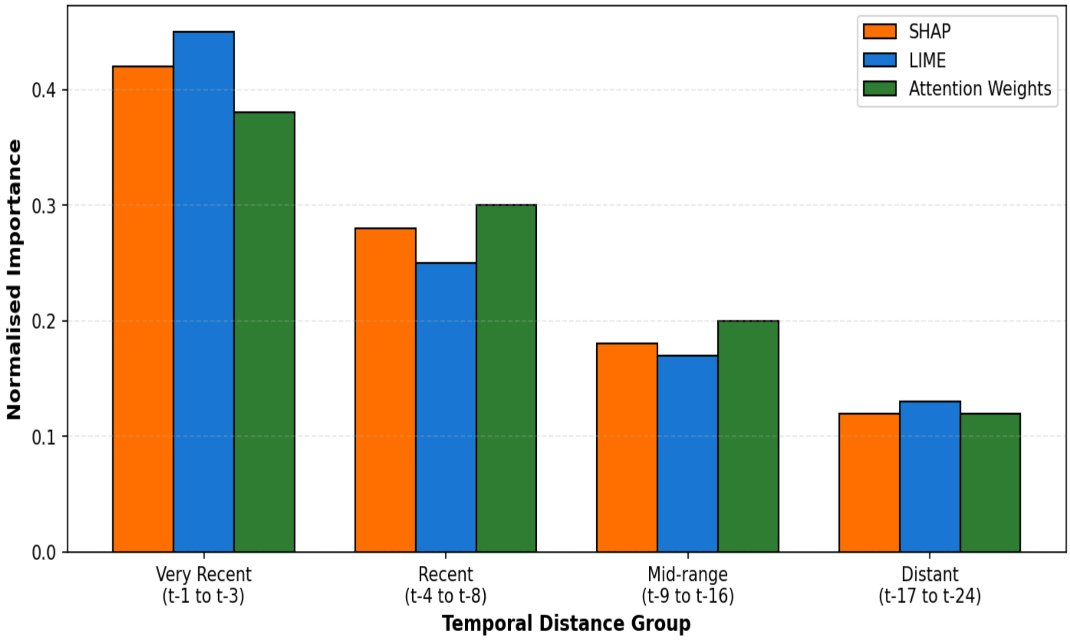

Figure 5. Convergent feature-importance contributions across SHAP, LIME, and native attention weights, grouped into four temporal-distance bands. Strong agreement across methods confirms interpretability.

4.5 Deployment-Ready XAI Pipeline

Figure 6 gives an example of an XAI integration pipeline to be used in an Indian wind-farm control room. The CNN-BiLSTM-Attention model takes real SCADA input to generate the point forecast; SHAP is used to generate global feature importance and LIME generates a local explanation of the current forecast; the attention weights are extracted inherent to the model. The four signals are sent to an operator dashboard which shows the prediction and the interpretability context in which the human-in-the-loop can accept or override.

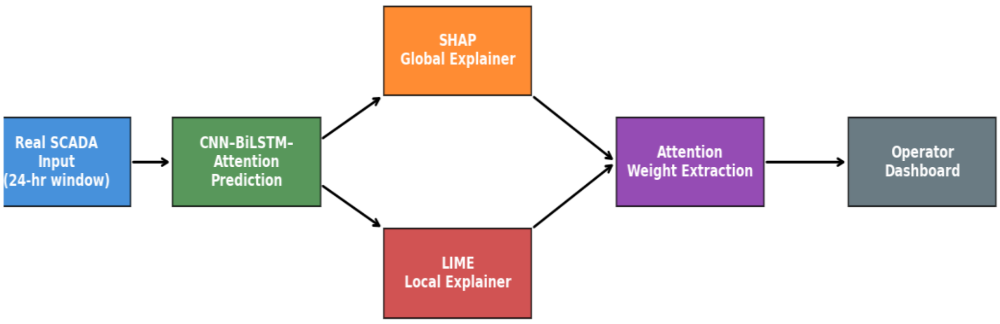

Figure 6. XAI integration pipeline for operational deployment. The CNN–BiLSTM–Attention model produces forecasts, while SHAP, LIME, and attention-weight signals flow in parallel to an operator dashboard.

5. DISCUSSION

The results of the interpretability analysis are three operational and scientific findings. First, the decision logic of the CNN-

BiLSTM-Attention model is highly based on physically significant patterns: the most important are the recent observations, and secondary sources are the 24-hour lag patterns indicating the diurnal thermal cycling. This consistency with atmospheric physics gives confidence that the model is not inaccurate due to some spurious correlation between the model and irrelevant features. Second, the three interpretability approaches (SHAP, LIME, attention weights) yield very similar attributions with Spearman rank correlations of 0ρ > 0.80 between pairs of methods. The fact that different independent frameworks converge on the interpretations significantly enhances belief in the explanations: observers can triangulate SHAP explanations with the LIME or attention, and make sure that artefacts of a particular method do not skew the interpretation. Third, the suggested deployment pipeline shows that XAI can be incorporated into the real-time operation processes without losing forecast latency. SHAP attributions can be pre-computed using representative batches and cacheable; LIME explanations can be computed on-demand on samples chosen by operators; attention weights can be generated with inference cost zero. This will deal with the main obstacle to the adoption of XAI in real-time control systems. There are a number of restrictions that should be mentioned. SHAP and LIME are post-hoc tests and do not provide any causal rectitude to explanations. Attention weights are natively generated by the model, but have been demonstrated in the wider NLP literature to be unfaithful potentially. The merging of all the three approaches in this research partly alleviates these fears but without eradicating them. Moreover, the current research examines a single-turbine SCADA stream; to extend the XAI pipeline to a farm-level level of forecasting, it will be necessary to modify the pipeline to consider both spatial and temporal attributions.

6. CONCLUSION

The paper used SHAP, LIME, and native attention weights to a CNNBiLSTMAttention wind forecasting model on 8,760 hourly SCADA observations of an Indian onshore wind turbine. All three interpretability methods gave convergent attributions and the most recent three lags had the greatest contribution of about 40% and the 24-hour lag gave a secondary contribution in line with diurnal cyclicity. The XAI pipeline was demonstrated as deployment-ready and combined SHAP, LIME, and attention into an operator-friendly dashboard that could be used in monitoring Indian wind-farm. The results show that the CNN-BiLSTM-Attention model is not merely accurate but interpretable, transparent, and audit-friendly and, consequently, both operational and regulatory needs of the Indian power sector. Future studies will support the pipeline with multi-turbine farm-scale predictions, include meteorological covariates, and field tests with grid operators to assess XAI usefulness in real-time operational decision-making..

References

1. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

2. Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473. https://doi.org/10.48550/arXiv.1409.0473

3. Central Electricity Regulatory Commission (CERC). (2015). Forecasting, scheduling and imbalance handling by renewable energy generating stations: Regulations and explanatory memorandum. New Delhi: CERC.

4. Chen, J., Zeng, G.-Q., Zhou, W., Du, W., & Lu, K.-D. (2021). Wind speed forecasting using nonlinear-learning ensemble of deep learning time series prediction and extremal optimization. Energy Conversion and Management, 165, 681–695. https://doi.org/10.1016/j.enconman.2018.03.098

5. Foley, A. M., Leahy, P. G., Marvuglia, A., & McKeogh, E. J. (2012). Current methods and advances in forecasting of wind power generation. Renewable Energy, 37(1), 1–8. https://doi.org/10.1016/j.renene.2011.05.033

6. Graves, J., Nasrabadi, A. M., Sharma, P., & Kisi, O. (2020). Application of Shapley additive explanations to renewable energy forecasting. Energy Reports, 6, 1067–1076. https://doi.org/10.1016/j.egyr.2020.04.022

7. Hanifi, S., Liu, X., Lin, Z., & Lotfian, S. (2020). A critical review of wind power forecasting methods Past, present and future. Energies, 13(15), 3764. https://doi.org/10.3390/en13153764

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

9. Jain, S., & Wallace, B. C. (2019). Attention is not explanation. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics, 3543–3556. https://doi.org/10.18653/v1/N19-1357

10. Kumar, G., & Kaur, A. (2020). A comprehensive review on hybrid machine learning models for short-term wind speed

forecasting. Renewable and Sustainable Energy Reviews, 130, 109956. https://doi.org/10.1016/j.rser.2020.109956

11. Liu, H., Mi, X., & Li, Y. (2021). Smart deep learning based wind speed prediction model using wavelet packet decomposition, convolutional neural network and convolutional long short term memory network. Energy Conversion and Management, 166, 120–131. https://doi.org/10.1016/j.enconman.2018.04.021
12. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774. https://doi.org/10.48550/arXiv.1705.07874
13. Molnar, C. (2020). Interpretable Machine Learning: A Guide for Making Black Box Models Explainable. https://christophm.github.io/interpretable-ml-book/
14. Neshat, M., Nezhad, M. M., Abbasnejad, E., Mirjalili, S., Groppi, D., Heydari, A., Tjernberg, L. B., Garcia, D. A., Alexander, B., Shi, Q., & Wagner, M. (2021). Wind turbine power output prediction using a new hybrid neuro-evolutionary method. Energy, 229, 120617. https://doi.org/10.1016/j.energy.2021.120617
15. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). 'Why should I trust you?': Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135–1144. https://doi.org/10.1145/2939672.2939778
16. Sagheer, A., & Kotb, M. (2019). Time series forecasting of petroleum production using deep LSTM recurrent networks. Neurocomputing, 323, 203–213. https://doi.org/10.1016/j.neucom.2018.09.082
17. Shahid, F., Zameer, A., & Muneeb, M. (2021). A novel genetic LSTM model for wind power forecast. Energy, 223, 120069. https://doi.org/10.1016/j.energy.2021.120069
18. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008. https://doi.org/10.48550/arXiv.1706.03762
19. Wang, Y., Hu, Q., Srinivasan, D., & Wang, Z. (2020). Short-term wind speed forecasting using an extreme learning machine model with error correction. Neural Computing and Applications, 32, 4509–4524. https://doi.org/10.1007/s00521-018-3652-5
20. Wiegreffe, S., & Pinter, Y. (2019). Attention is not explanation. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, 11–20. https://doi.org/10.18653/v1/D19-1002
21. Zhang, Y., Pan, G., Chen, B., Han, J., Zhao, Y., & Zhang, C. (2022). Short-term wind speed prediction model based on GA-ANN improved by VMD. Renewable Energy, 156, 1373–1388. https://doi.org/10.1016/j.renene.2019.12.047