Annual Peak Load Forecasting for Distribution Feeders Using AMI Data and LSTM

Jongchan Choi, Jin Dong, Ben Ollis, Jamie Lian
Oak Ridge National Laboratory
Oak Ridge, USA
choij1, dongj, ollistb, lianj@ornl.gov

Samuel Levinson
Duquesne Light Company
Pittsburgh, USA
slevinson@duqlight.com

Jeremy Keen and Thomas Bowen
National Laboratory of the Rockies
Golden, USA
Jeremy.Keen, Thomas.Bowen@nrel.gov

***Abstract*—Accurate peak load forecasting is critical for the reliable operation and strategic planning of distribution systems. This paper proposes a Long Short-Term Memory (LSTM)-based framework that integrates advanced metering infrastructure (AMI) data with key weather indicators—including heating degree days, cooling degree days, extreme hot days, extreme cold days, persistent hot days, and persistent cold days—for annual peak load forecasting. Unlike conventional approaches, the proposed method provides enhanced spatial and temporal granularity, enabling peak forecasts not only at the system level but also at individual feeders and fine time resolutions. To mitigate systematic bias, historical forecasting errors are systematically analyzed and incorporated to model recalibration. The proposed framework is evaluated at both system and feeder levels, achieving forecasting errors consistently within 4%, which significantly outperforms conventional approaches.**

***Index Terms*—Advanced metering infrastructure (AMI), load forecasting, peak load, distribution feeder, long short-term memory (LSTM), weather indicators, recalibration.**

I. Introduction

Accurate annual peak load forecasting is fundamental to distribution network planning, enabling timely investments in transformers, substations, feeders, advanced metering infrastructure (AMI), and other smart-grid assets. Rapid changes driven by distributed energy resources (DERs), data-center growth, electrification, evolving consumption with demand response, and extreme weather heighten forecasting challenges, requiring frequent recalibration [1], sensitivity to weather [2]–[4], and feeder-level complexity [5], [6]. Moreover, with U.S. residential electricity prices projected to rise steadily—13% in 2025 and 18% in 2026 relative to 2022 [7], outpacing inflation and other retail fuels—improved forecasting is critical to reduce costly over- or under-investment, mitigate upward rate pressure, and enhance long-term affordability. In addition, distribution investments have emerged as a key driver of total U.S. utility spending, steadily increasing over the past two decades and now representing a growing share of overall capital investment in electricity infrastructure [8].

This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725, and NREL under contract DE-AC36-08GO28308 with the US Department of Energy (DOE). Funding provided by US Department of Energy Office of Energy Efficiency and Renewable Energy, Energy Grid Integration Office.

State-of-the-art annual forecasting for distribution planning combines traditional statistical and deterministic methods, scenario analysis, and advanced computational techniques. Utilities typically rely on historical load profiles, billing records, demographic and socioeconomic indicators, and land-use data to form baseline projections. Hybrid strategies—especially ensemble approaches and stacking frameworks—routinely outperform single-method solutions by improving robustness to model drift and weather-driven uncertainty [9], [10]. Advanced stacking that fuses diverse artificial intelligence (AI) models with meta-learners enhances accuracy for capturing multivariate dependencies and nonlinear load dynamics [11].

Rigorous uncertainty quantification is increasingly essential for risk-aware decision-making. Recent developments include diffusion-based methods that separate epistemic and aleatoric uncertainty and probabilistic forecasting techniques that produce confidence intervals instead of point estimates [12], [13]. These approaches help planners account for dynamic and uncertain consumption patterns. Advances in AI—transformers, attention mechanisms, and neural architecture search—have further improved the modeling of temporal and spatial dependencies, with time-pattern attention improving feature extraction for load prediction [12]–[16].

AMI data helps utilities operate more effectively in a variety of ways including improving outage management, enhancing billing accuracy, supporting demand response programs, enabling grid optimization, and strengthening load forecasting for infrastructure planning [17]. AMI meters generate data at intervals of 15 minutes or less, enabling more accurate load modeling and improved grid planning [18]–[20]. However, leveraging AMI for annual peak load forecasting remains challenging due to the volume and complexity of the data, as well as ongoing concerns related to data quality, privacy, and integration with forecasting models.

Despite recent advances, annual peak load forecasting remains challenged by emerging large loads, electrification, and evolving consumer behavior. Climate variability intensifies extreme weather, driving unpredictable record peaks. Existing models often deteriorate over time, offer limited uncertainty quantification, and lack feeder-level precision. To address these

979-8-3315-8138-1

gaps, we propose a Long Short-Term Memory (LSTM)-based framework that leverages AMI data, weather indicators, and historical error recalibration to enhance robustness, accuracy, and interpretability, validated on real-world field data.

II. Annual Peak Load Forecasting for Distribution Grids

A. Current Load Forecasting Practices at the Utility

The objective of the annual peak load forecast is to develop an accurate projection for each distribution circuit based on historical loading patterns, economic development trends, and load growth behavior. In practice, distribution utilities combine data sources such as historical load profiles, customer billing records, demographic information, land-use and zoning plans, and localized economic indicators to establish a reliable baseline forecast. Practices, however, vary widely; it is common for some utilities to rely on simpler approaches, such as weather-normalized trends of past loading, particularly when data availability is limited. The study is conducted in collaboration with a regional utility company within the Pennsylvania–New Jersey–Maryland (PJM) Interconnection area, based on real-world operational practices.

For the first forecast year, all distribution circuit's peak load is estimated using a load growth factor (LGF) for the utility's service territory published in the PJM Load Forecast Annual Report [21]. This initial estimate is adjusted to account for new load connections expected within the year as well as anticipated load transfers from other circuits. The calculation follows:

$$P_{\text{next}} = (L_{\text{new}} + P_{\text{last}}) \times (1 + LGF_{PJM}) \tag{1}$$

where, P_{next} is the next year's projected peak load, L_{new} is the newly added or transferred load connections to the system, P_{last} is the previous year's peak, and LGF_{PJM} is the growth factor expressed as a decimal (e.g., 0.02 for 2%).

From the second forecasted year onward, a fixed LGF of 1.5% is applied, based on the utility's historical data, load growth trends, and customer class segmentation. This long-term projection uses the following equation:

$$P_{\text{next}} = P_{\text{last}} \times (1 + LGF_{historical}) \tag{2}$$

where, $LGF_{historical}$ represents the fixed 1.5% LGF.

Once the forecast is generated by (1) and (2), the results are compared against circuit ratings to identify two key conditions: (i) overload risks, where the projected peak load exceeds circuit capacity, and (ii) high-load conditions, where the peak load reaches or exceeds 90% of capacity. Overload risks are then cross-checked against planned system improvement projects. If the issue is already addressed by ongoing or proposed work, no additional intervention is required. If not, new projects are initiated—often involving load transfers to adjacent circuits, conductor upgrades, or substation expansions. Peak load determinations incorporate detailed operational data from substation breakers and fault interrupters, adjusted for any scheduled load transfers.

While effective for broad planning purposes, the current forecasting approach has three major limitations. First, the approach treats all feeders uniformly, potentially masking local unique growth patterns, constraints, or operational behaviors at the individual feeder level. Second, the method does not explicitly integrate weather normalization or account for extreme temperature events, both of which can significantly influence annual peaks. Third, reliance on a fixed load growth factor may overlook demand variations and emerging trends such as electrification or DER growth. Addressing these limitations through more granular, dynamic, and data-driven methods could significantly improve forecast accuracy and planning efficiency.

B. AMI data and Load Forecasting

Utilities are increasingly leveraging AMI data to enhance operational efficiency and gain deeper insights into grid performance for improved situational awareness. However, applying AMI data to load forecasting remains complex due to its massive volume, high dimensionality, and computational demands. Data from millions of meters require advanced analytical methods and substantial processing resources. In addition, household-level consumption is inherently volatile and difficult to predict, while existing forecasting models—typically built on aggregated, system-level data—are not readily adaptable to such granular inputs. Integrating AMI data into these legacy frameworks is therefore both challenging and resource-intensive. Nonetheless, as data analytics and modeling techniques advance, AMI is poised to become a cornerstone for more accurate, scalable, and distribution-level load forecasting.

III. Annual Peak Load Forecasting by LSTM

A. LSTM Networks

LSTM networks are well-suited for load forecasting using load profiles because they can effectively capture temporal dependencies and nonlinear dynamics within sequential data [22]. By analyzing the sequence of daily peaks across a year, LSTMs can learn how short-term events, seasonal patterns, and irregular spikes contribute to the annual maximum for annual peak load forecasting. This approach leverages the richness of daily data, improves robustness to anomalies, and enables the model to account for the complex factors that influence annual peak demand.

An LSTM cell maintains a cell state (C_t) and a hidden state (h_t), with information flow regulated by three gates: the forget gate, the input gate, and the output gate. The architecture is shown in Fig. 1. The mathematical representation of an LSTM cell is:

$$f_t = \sigma(W_f \cdot [h_{t-1}, \mathbf{X}_t] + b_f) \tag{3}$$
$$i_t = \sigma(W_i \cdot [h_{t-1}, \mathbf{X}_t] + b_i) \tag{4}$$
$$\tilde{C}t = \tanh(W_c \cdot [ht-1, \mathbf{X}_t] + b_c) \tag{5}$$
$$C_t = f_t * C_{t-1} + i_t \odot \tilde{C}t \tag{6}$$
$$o_t = \sigma(W_o \cdot [h_{t-1}, \mathbf{X}_t] + b_o) \tag{7}$$
$$h_t = o_t \odot \tanh(C_t) \tag{8}$$

979-8-3315-8138-1

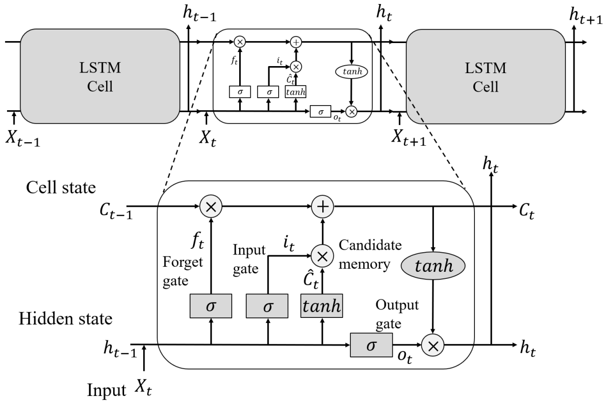

Fig. 1. LSTM architecture and cell structure

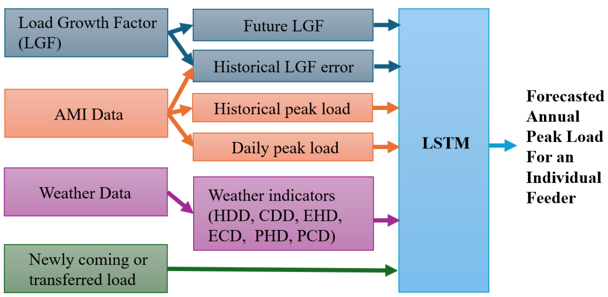

Fig. 2. Inputs and output of the proposed LSTM framework

where X_t is the input at time t, h_{t-1} and C_{t-1} are the hidden and cell states from the previous time step, σ is the sigmoid activation function, W and b are trainable weights and biases, and $*$ denotes element-wise multiplication.

In peak load forecasting, LSTM networks capture both seasonal patterns and irregular variations by learning from historical demand. They support point or multi-step forecasting, typically evaluated via Mean Squared Error (MSE).

To model extreme weather impacts, we employ a multivariate LSTM architecture. The input vector $\mathbf{X}_t$ explicitly incorporates historical load alongside the proposed intensity and persistence indicators.

This dynamic, data-driven approach replaces static growth assumptions, adapting to localized patterns, weather anomalies, and emerging trends. Consequently, LSTM-based models offer a more accurate and responsive alternative that strengthens traditional planning methods.

B. Inputs to LSTM Network for Annual Peak Load Forecasting

Multiple inputs including LGF, weather data, and AMI data are incorporated with the proposed LSTM framework for annual peak load forecasting as shown in Fig. 2

1) LGF and Historical Bias: The recalibration of LGFs, also referred to as bias correction, is particularly effective when sufficient historical data is available. The process begins by estimating feeder-level LGF errors through a comparison of observed historical values with those reported by PJM. These errors are evaluated at both the system-wide and feeder-specific levels. The comparison highlights that feeder-level recalibration yields substantial improvements in the accuracy of annual peak load forecasts as listed in Table I.

TABLE I
HISTORICAL LGF ERRORS FOR INDIVIDUAL FEEDER AND REGIONAL SYSTEM [%]

Year	Feeder 1	Feeder 2	Feeder 3	Feeder 4	Feeder 5	Regional Average
2019	17.97	7.52	19.98	8.81	4.40	14.78
2020	-4.97	3.64	1.74	6.69	-1.72	2.88
2021	-16.98	65.53	-5.13	16.96	-9.24	-0.42
2022	-1.18	-48.77	7.94	-2.71	30.36	6.04
2023	19.88	44.08	-0.77	-13.46	-14.28	9.45
2024	4.42	1.45	18.19	1.16	-19.43	3.29
Averaged	**10.90**	**28.50**	**8.96**	**8.30**	**13.24**	**6.00**

*positive: over-forecast, negative: under-forecast

The resulting feeder-specific LGF error is incorporated into the developed LSTM network as an input feature, alongside the PJM-reported LGF. This dual-input design enables the model to capture both systematic and localized biases, thereby enhancing the robustness and accuracy of load forecasting.

2) Expansion of Weather Feature Representation: Traditional peak load forecasting approaches often rely on aggregate seasonal indicators such as Heating Degree Days (HDD) and Cooling Degree Days (CDD).

- Heating Degree Days (HDDs): Measure the extent to which daily temperatures fall below a standard threshold (commonly 65°F or 18°C), serving as a primary indicator of heating-driven electricity demand.
- Cooling Degree Days (CDDs): Capture the extent to which daily temperatures exceed the same threshold (commonly 65°F or 18°C), serving as a direct proxy for cooling loads and summer grid stress.

While HDDs and CDDs provide useful seasonal benchmarks, they often fail to capture the short-term, high-intensity weather phenomena that drive peak demand. Higher seasonal totals generally imply increased overall consumption but do not reliably predict peak days, as peak events are disproportionately shaped by isolated major events rather than seasonal averages.

To address these limitations, this study expands weather feature representation to include explicit descriptors of intensity and duration. All indicators are calculated at a daily resolution.

- Extreme Cold Days (ECDs): Temperatures fall below a critical threshold (e.g., <10°F or -12°C, or below the 5th percentile of historical values).
- Extreme Hot Days (EHDs): Temperatures exceed a critical threshold (e.g., >85°F or 29.5°C, or above the 95th percentile).

They are quantified as binary flags (EHD, ECD $\in \{0, 1\}$). EHDs and ECDs take a value of 1 when the daily mean temperature exceeds the 95th percentile or falls below the 5th percentile of historical values, respectively, and 0 otherwise. These extreme days frequently coincide with peak load conditions, as demand rises sharply in response to sudden thermal change.

Another critical dimension is persistence, defined as the occurrence of multiple consecutive days of extreme or near-extreme temperatures:

- Persistent Cold Days (PCDs): Multi-day periods of unusually low temperatures that sustain elevated heating

979-8-3315-8138-1

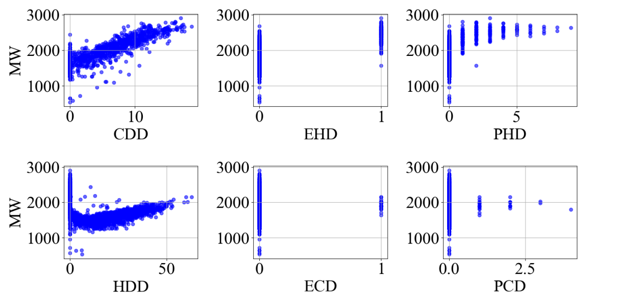

Fig. 3. Relationship between daily peak load and weather indicators including CDD, HDD, EHD, ECD, PHD, and PCD for 2019-2024

demand.

- Persistent Hot Days (PHDs): Multi-day periods of elevated temperatures that prevent cooling systems from recovering between cycles.

PCDs and PHDs are as count-based variables (unit: days). These metrics measure the cumulative duration of consecutive extreme days using a recursive logic: the counter increments for every subsequent day ECD or EHD continues (e.g., 0, 1, 2...), but resets to zero immediately if the sequence is broken.

These multi-day events impose cumulative stress on both end-use systems and grid infrastructure. Persistence introduces a compounding effect: continuous HVAC operation erodes building thermal inertia, while behavioral adaptations (e.g., thermostat overrides) amplify demand intensity. Consequently, persistent events often produce the highest system loads due to cumulative stress, even without record-breaking daily temperatures.

As shown in Fig. 3, CDD and HDD exhibit a strong linear correlation with load. In contrast, PHD and PCD show greater dispersion. This scatter is expected because persistence is not an independent driver but an amplifying factor. A specific persistence value (e.g., PHD=2) can coincide with varying absolute temperatures. Thus, its explanatory power is fully realized only when combined with intensity metrics, allowing the model to differentiate between isolated extremes and compounding multi-day stress.

3) Historical and Imminent Loads for Individual Feeder: An aggregated AMI data at the feeder level is leveraged to improve load forecasting, with particular focus on capturing historical peak demand and quantifying the effects of newly added or transferred loads. Data from individual AMI devices is aggregated for each feeder, as indicated by the blue circle markers in Fig. 4. Then the AMI data is integrated into the developed LSTM network as an input feature, alongside records of load additions and transfers. This combined representation enables the model to account for localized growth dynamics arising from demographic changes, electrification trends, and structural shifts in end-use consumption. By embedding feeder-level characteristics within the forecasting framework, the proposed approach improves both the robustness and accuracy of peak load predictions.

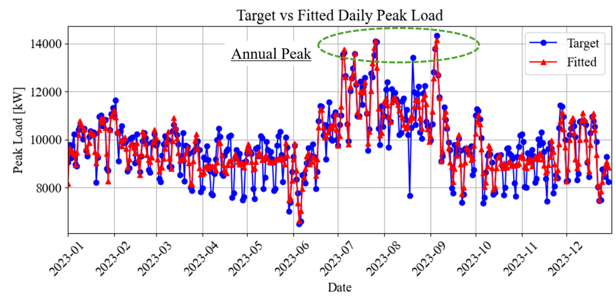

Fig. 4. A sample daily peak load profile (Target vs fitted curve)

IV. Case Study

A. Training and Validation

1) Data Preparation: The training process was carried out using historical load and weather data sets described in Section III, structured into a supervised learning format for the LSTM model. The dataset was first preprocessed by separating inputs and target (i.e., daily peak load profile) variables, followed by normalization to ensure stable convergence during training. Input features consisted of eleven variables, including historical peak load, LGF, LGF bias, weather indicators (CDD, HDD, EHD, ECD, PHD, and PCD), and temporal attributes, while the output variable corresponded to the daily peak load.

2) Sequence Construction and Batching: To effectively capture temporal dependencies, the input data were reshaped into sliding sequences of 90 days, ensuring that each training sequence represented a full seasonal cycle. Both the input and target sequences were stacked into tensors to be used with PyTorch's Dataset and DataLoader modules, allowing efficient batching and shuffling during training.

3) Model Architecture: The forecasting model was implemented as a stacked LSTM network with two recurrent layers. Each layer contained 64 hidden units, followed by a fully connected layer that mapped the hidden representations to the output space. Dropout regularization was incorporated in the recurrent layers to mitigate overfitting, while the final dense layer provided sequence-to-sequence predictions of peak loads.

4) Training Configuration: Training was performed using the Adam optimizer with an initial learning rate of 1e-3, and MSE was used as the loss function. The model was trained for 500 epochs on a CPU, and average training loss per epoch was recorded for performance monitoring.

5) Validation Strategy: The validation strategy focused on monitoring training loss convergence and ensuring that the model generalized well across sequences. By leveraging normalized features, long-sequence training, and error recalibration mechanisms, the LSTM model was able to learn robust temporal dependencies while maintaining forecasting accuracy within the desired error bounds, as indicated by the red triangular markers in Fig. 4.

6) Scalability and Implementation: While this study evaluates the framework on a representative set of feeders, the proposed architecture is designed to scale to larger distribution systems. The data-driven nature of the LSTM model allows for parallelized training and inference, where individual models

979-8-3315-8138-1

TABLE II
ERRORS FOR 2024 ANNUAL PEAK LOAD FORECASTING FOR SAMPLE FEEDERS [%]

	Feeder 1	Feeder 2	Feeder 3	Feeder 4	Feeder 5	Averaged
W/ Current approach	19.58	46.3	29.64	15.97	-11.12	20.07
W/ LSTM Only	-30.7	13.1	-22.7	-40.5	-19.3	-20.0
W/ Proposed LSTM with AMI/weather	1.20	-4.02	-0.72	1.25	-0.27	-0.51

*positive: over-forecast, negative: under-forecast

can be trained for specific feeders or clusters of feeders using distributed computing resources.

B. Case Study: 2024 Peak Load Forecasting

The forecasting performance of the proposed approach was evaluated using annual peak load data for 2024 across five representative feeders. Absolute error was selected as the evaluation metric, with results benchmarked against the current utility forecasting approach described in Section II. As reported in Table II, the conventional method exhibited errors ranging from -11.12% (under-forecast) to 46.3% (over-forecast), reflecting substantial variability and limited reliability at the feeder level. In contrast, the proposed LSTM-based approach, which incorporates AMI data, consistently achieved errors below 4%, with the minimum recorded at 1.13%. This represents a significant improvement in predictive accuracy, underscoring the robustness and adaptability of the proposed model. In addition, by limiting errors to within 4%, the proposed model enhances grid reliability, and contributes to cost savings in operational planning.

V. CONCLUSION

This paper presents a data-driven LSTM-based framework for annual peak load forecasting that integrates high-resolution AMI data with advanced weather indicators, including measures for extreme intensity and persistence. By incorporating a historical error recalibration mechanism, the framework effectively mitigates systematic bias and adapts to localized feeder characteristics. Validation on real-world data demonstrates that the proposed approach consistently restricts forecasting errors to within 4%, significantly outperforming conventional utility methods which exhibited errors as high as 46%. These results confirm that granular, data-driven forecasting can reduce reliance on over-provisioned capacity, enabling more efficient infrastructure planning and enhanced grid reliability under increasingly variable weather conditions.

ACKNOWLEDGMENT

The project is funded by the Energy Grid Integration Office within the Office of Energy Efficiency and Renewable Energy.

REFERENCES

[1] A. Azeem, I. Ismail, S. M. Jameel, F. Romlie, K. U. Danyaro, and S. Shukla, "Deterioration of Electrical Load Forecasting Models in a Smart Grid Environment," *Sensors*, vol. 22, no. 12, 2022.

[2] D. Cawthorne, A. R. de Queiroz, H. Eshraghi, A. Sankarasubramanian, and J. F. DeCarolis, "The Role of Temperature Variability on Seasonal Electricity Demand in the Southern US," *Frontiers in Sustainable Cities*, vol. Volume 3, 2021.

[3] New York Independent System Operator (ISO), "Load Forecast Uncertainty Modeling: New York Temperature Distribution Analysis," New York ISO, Tech. Rep., 2021.

[4] G. Liu, J. Liu, Y. Bai *et al.*, "EWELD: A large-scale industrial and commercial load dataset in extreme weather events," *Scientific Data*, vol. 10, p. 615, 2023.

[5] A. Aditya, K. S. Gahlot, and S. Joshi, "Hierarchical BiLSTM-LSTM Network for Short-Term Load Forecasting," in *2024 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pp. 1–3, 2024.

[6] M. Grabner, Y. Wang, Q. Wen, B. Blažič, and V. Štruc, "A Global Modeling Framework for Load Forecasting in Distribution Networks," *IEEE Transactions on Smart Grid*, vol. 14, no. 6, pp. 4927–4941, 2023.

[7] U.S. Energy Information Administration, "U.s. electricity prices continue steady increase," https://www.eia.gov/todayinenergy/detail.php?id=65284, 2025, accessed: 2025-10-01.

[8] U.S. Energy Information Administration, "Grid infrastructure investments drive increase in utility spending over last two decades," https://www.eia.gov/todayinenergy/detail.php?id=63724, 2024, accessed: 2025-10-13.

[9] H. Xi, W. Senyao, L. Ni, P. Junyang, L. Yixian, L. Jingfeng, and W. Xueting, "Ensemble Learning for Short-Term Load Forecasting: Integrated Methods of Multiple Individual Models," in *2024 CEECT*, pp. 1–5, 2024.

[10] K. R. Hossain, M. S. S. Islam, M. R. Islam, M. R. I. Sheikh, and M. O. Faruque, "CNN-XGBoost Based Ensemble Model for Short-Term Load Forecasting," in *2024 IEEE RAAICON*, pp. 30–33, 2024.

[11] F. Guo, H. Mo, J. Wu, L. Pan, H. Zhou, Z. Zhang, L. Li, and F. Huang, "A hybrid stacking model for enhanced short-term load forecasting," *Electronics*, vol. 13, no. 14, 2024.

[12] Z. Wang, Q. Wen, C. Zhang, L. Sun, and Y. Wang, "DiffLoad: Uncertainty Quantification in Electrical Load Forecasting with the Diffusion Model," 2024. [Online]. Available: https://arxiv.org/abs/2306.01001

[13] V. Álvarez, S. Mazuelas, and J. A. Lozano, "Probabilistic Load Forecasting Based on Adaptive Online Learning," *IEEE Transactions on Power Systems*, vol. 36, no. 4, pp. 3668–3680, 2021.

[14] B. Jiang, H. Yang, Y. Wang, Y. Liu, H. Geng, H. Zeng, and J. Ding, "Dynamic Temporal Dependency Model for Multiple Steps Ahead Short-Term Load Forecasting of Power System," *IEEE Transactions on Industry Applications*, vol. 60, no. 4, pp. 5244–5254, 2024.

[15] W. Liao, J. Ruan, Y. Xie, Q. Wang, J. Li, R. Wang, and J. Zhao, "Deep learning time pattern attention mechanism-based short-term load forecasting method," *Frontiers in Energy Research*, vol. Volume 11, 2023.

[16] X. Wen, J. Liao, Q. Niu, N. Shen, and Y. Bao, "Deep learning-driven hybrid model for short-term load forecasting and smart grid information management," *Scientific Reports*, vol. 14, no. 1, p. 13720, 2024.

[17] K. Zhou, C. Fu, and S. Yang, "Big Data Driven Smart Energy Management: From Big Data to Big Insights," *Renewable and Sustainable Energy Reviews*, vol. 56, pp. 215–225, 2016.

[18] J. Xu, M. Yue, D. Katramatos, and S. Yoo, "Spatial-temporal load forecasting using ami data," in *2016 IEEE SmartGridComm*, 2016.

[19] M. H. Leak and G. K. Venayagamoorthy, "Forecasting peak daily load in distribution feeders," in *2018 Clemson University Power Systems Conference (PSC)*, 2018.

[20] A. Khazaee, H. Chenarani, H. Delavaripour, and R. Afshar, "Smart planning of distribution network using ami system," in *2023 27th International Electrical Power Distribution Networks Conference (EPDC)*, 2023.

[21] "PJM Long-Term Load Forecast Report," Pennsylvania-New Jersey-Maryland Interconnection (PJM), Valley Forge, PA, Tech. Rep., Jan. 2025.

[22] W. Kong, Z. Y. Dong, Y. Jia, D. J. Hill, Y. Xu, and Y. Zhang, "Short-term residential load forecasting based on lstm recurrent neural network," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 841–851, 2019.