

Cyclical Temporal Encoding and Hybrid Deep Ensembles for Multistep Energy Forecasting

Salim Khazem^{*1} Houssam Kanso¹

Abstract—Accurate electricity consumption forecasting is essential for demand management and smart grid operations. This paper introduces a unified deep learning framework that integrates cyclical temporal encoding with hybrid LSTM–CNN architectures to enhance multistep energy forecasting. We systematically transform calendar-based attributes using sine–cosine encodings to preserve periodic structure, and evaluate their predictive relevance through correlation analysis. To exploit both long-term seasonal effects and short-term local patterns, we employ an ensemble model composed of an LSTM, a CNN, and a meta-learner of MLP regressors specialized for each forecast horizon. Using a one-year national consumption dataset, we conduct an extensive experimental study including ablation analyses (with/without cyclical encodings and calendar features) and comparisons with established baselines from the literature. Results demonstrate consistent improvements across all seven forecast horizons, with our hybrid model achieving lower RMSE and MAE than individual architectures and prior methods. These findings confirm the benefit of combining cyclical temporal representations with complementary deep learning structures. To our knowledge, this is the first work to jointly evaluate temporal encodings, calendar-based features, and hybrid ensemble architectures within a unified short-term energy forecasting framework. The code is available at <https://github.com/salimkhazem/cai2025.git>

I. INTRODUCTION

Accurate forecasting of electricity consumption remains a critical challenge for energy providers and smart grid operators. It is essential for ensuring the balance between supply and demand, optimizing energy storage systems, and informing strategic decisions such as dynamic pricing and load management. However, electricity usage is shaped by a multitude of complex and interdependent factors including human behavior, seasonal climate variability, and socio-economic dynamics that introduce strong non-linearities and periodic patterns into consumption trends. These intricacies often limit the effectiveness of traditional statistical forecasting approaches [1].

Recent advances in deep learning have shown substantial promise in tackling time series forecasting problems. Notably, architectures such as Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNNs) are capable of learning long-range temporal dependencies and extracting meaningful patterns from historical sequences [2]. Yet, their performance remains highly sensitive to the way temporal information is represented. Raw time step encoding alone often fails to capture the underlying periodic structure inherent in electricity consumption

data. Hybrid and ensemble architectures have shown strong potential in load forecasting due to their ability to model both local and long-range temporal patterns [3], [4].

In this work, we propose an enhanced deep learning-based forecasting framework that incorporates explicit temporal feature engineering. We extract and analyze calendar-based attributes such as day-of-week, week-of-year, and month-of-year, which are known to exhibit strong cyclical influence on consumption behaviors. To preserve their periodic nature and avoid discontinuities (e.g., transitions from Sunday to Monday or December to January), we apply sine and cosine transformations. These transformations map discrete time features onto a continuous cyclic space, enabling neural networks to learn seasonal rhythms more effectively.

We conduct a correlation-driven feature selection process to assess the predictive power of each temporal attribute, in both raw and transformed formats. Features with the strongest correlation to the target variable are incorporated into the input space of the forecasting models. We then evaluate the performance of three architectures: LSTM, CNN, and a hybrid model over a rolling seven-day prediction horizon, using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) as evaluation metrics [5].

Our results demonstrate that the integration of well-encoded temporal features significantly enhances forecasting accuracy, particularly for medium and long-term horizons. This study contributes a robust and interpretable temporal encoding pipeline that can be generalized to other periodic time series in smart energy systems and beyond.

Our work addresses two limitations observed in prior research. First, although cyclical time encodings (sine–cosine transforms) are widely used in traffic and energy forecasting, prior studies rarely quantify when these encodings are beneficial or how strongly each temporal feature correlates with the target signal. Second, hybrid deep learning architectures combining CNN and LSTM have been explored, but existing works do not examine how temporal encodings and calendar features interact with ensemble predictors across multi-day horizons.

To fill these gaps, our contribution is threefold:

- A systematic analysis of raw vs. cyclically transformed temporal features, using Pearson correlation to quantify how well each feature aligns with real consumption dynamics. This offers a principled justification for when cyclical encoding improves predictability, an aspect not explicitly studied in prior load forecasting literature.
- A unified hybrid ensemble architecture combining LSTM, CNN, and horizon-specific MLP meta-learners.

^{*}Corresponding author

¹Salim Khazem & Houssam Kanso are with Talan Research Center, Paris, France salim.khazem@talan.com

Unlike prior CNN–LSTM hybrids that fuse representations at a single level, our architecture jointly models long-term seasonality, local weekly structure, and day specific error refinement.

- A comprehensive evaluation, including: (i) ablation studies isolating the effects of cyclical encoding and calendar features, (ii) comparisons against representative state-of-the-art models, and (iii) analysis across a 7-day rolling horizon on a real national load dataset.

Together, these contributions provide new insights into the role of temporal encoding in multistep energy forecasting and demonstrate that hybrid–ensemble architectures combined with well-designed time features yield consistent accuracy gains.

II. RELATED WORK

Electric load forecasting has long been a critical area of research, particularly with the advent of smart metering technologies and the increasing emphasis on demand-side management in smart grids. Early approaches have predominantly relied on statistical models, such as autoregressive methods like ARIMA [6], which assume linearity and stationarity. More recently, machine learning and deep learning techniques have emerged as powerful alternatives for capturing non-linear dependencies and complex temporal dynamics [2], [7].

Recent surveys further highlight the effectiveness of deep neural networks for short-term and multistep load forecasting [2], [8]. CNN-based architectures have also been leveraged to capture local patterns and short-term fluctuations in electricity demand [9]. Hybrid CNN–LSTM architectures have shown strong performance by combining sequential modeling with convolutional feature extraction [10], [11]. These architectures have also been applied successfully to capture multi-scale temporal structures [9], [12].

A growing body of work highlights the importance of temporal feature engineering in enhancing forecasting performance. Deep learning models, including LSTM [13] and CNN [14] architectures, have shown significant improvements when augmented with calendar based features [15]. To better represent the cyclical nature of time variables such as day-of-week, hour-of-day, or month-of-year, several studies adopt sine and cosine transformations [16]. These encodings help mitigate artificial discontinuities in periodic signals, thereby improving model generalization. Such techniques have been successfully applied in domains like traffic prediction [17] and energy demand forecasting [5].

Additionally, some approaches incorporate exogenous variables, such as temperature, holidays, and socio-economic factors, to further refine predictions [18]. These variables often introduce important contextual signals that influence consumption patterns beyond time alone.

III. METHOD

A. Preprocessing & Feature Engineering

1) **Exploration of Time data granularity:** Electric power consumption is split into several activity fields belonging

to domestic and professional uses. Those activity fields do not follow the same frequency of power consumption. In addition, several aspects affect the power utilization such as human behaviour, climatic conditions, and seasonal variation. For this purpose, we proceed to extract all time granularities in order to explore their relation with respect to power consumption. We extract from a single date the features of Day in Week, Day in Month, Day in Year, Week in Month, Week in Year, Month in Year. In addition, we generate an additional feature modeling the “business day” including the French national days off.

Then, we study the evolution of each feature with respect to the power consumption, by generating the Pearson correlation coefficient r , which is computed as indicated in the equation.

$$r = \frac{\sum (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum (x_i - \bar{x})^2} \cdot \sqrt{\sum (y_i - \bar{y})^2}}$$

Where x_i is the feature, y_i is the power consumption, $\bar{x}$ and $\bar{y}$ are respectively the feature and the power consumption average. The r coefficient models a linear relationship between the provided data, where a common rise (e.i. reversed rise) leads to a positive coefficient value (e.i. negative coefficient value). All r coefficients with respect to the time features are indicated in Table I.

TABLE I: Correlation Coefficients of Time Features

Feature	Correlation (r)
Day in Week	-0.1993
Day in Month	0.0370
Day in Year	-0.2338
Week in Month	-0.0022
Week in Year	-0.2288
Month in Year	-0.2341
Business Day	-0.2733

We note that *Day in Month* and *Week in Month*, where values are respectively in [1, 31] and [1, 6], have no correlation with the power consumption and are hence ignored for the rest of the study.

We explore the feature tendency in terms of power consumption in order to detail the r coefficients. The power variation in terms of months and weeks are respectively illustrated in boxplot visuals of Fig. 1a and Fig. 1b. We deduce a cyclic pattern of power consumption along an entire year, which explains the significant r coefficient related to “Day in Year”, “Week in Year” and “Month in Year”. A similar exploration has been done in terms of “Day in Week”, where we note an almost constant power consumption at the beginning of the week, as indicated by the median values in the boxplot of Fig. 1c and a power consumption decrease during the weekend. The same pattern is also deduced from Fig. 1d, modelling power in terms of “Business Day”.

2) **Feature Engineering for Cyclical Pattern of Time Series:** Cyclical patterns along the weeks and along the years are well modeled either in the whole dataset or in visual distributions. However, the related r coefficients are not high

enough to model such correlation. In fact, while power consumption evolves continually through the cyclical patterns, time features are characterized by a value discontinuity, due to value reset at the end of each cycle, as will distinguish in 2.

To overcome this problem data cut-off, we suggest transforming the timing data feature in order to have a progressive evolution along the time pattern. Within this objective, two news features are generated from each original ones, that correspond respectively to the sin and cos processing with respect to the identified DataCycle, as indicated in Eq. 1 and Eq. 2.

$$F_{sin} = \sin \left(2\pi \times \frac{F}{TimeCycle} \right) \quad (1)$$

$$F_{cos} = \cos \left(2\pi \times \frac{F}{TimeCycle} \right) \quad (2)$$

Where F is the time feature and $TimeCycle$ in the cycle duration of F . Figures in 2 illustrate the transformation of the “Week in Year” feature as indicated in equations 1 and 2, where trigonometry data evolutions are closed to the power consumption.

To evaluate the contribution made by these transformation, we generate the r coefficient for the original and the transformed features, as depicted in tab.(3). We note that F_{sin} and F_{cos} have higher correlation with the power consumption as well as preserving time cyclical time pattern. The Business Day remains unchanged even with trigonometric transformation. The “Day in Week” has a similar correlation after applying eq. 1, due to the consumption spike between Sunday and Monday, which closed to the evolution of day numbers. Consequently, we choose the transformed features offering the higher r coefficients for “Day in Year”, “Week in Year” and “Month in Year”, as modeled in bold in table (2). The “Day in week” and “Business Day” are preserved without transformation.

TABLE II: Correlation Coefficients in Terms of Trigonometry Transformations

	Day in Year	Week in Year	Month in Year	Day in Week	Business Day
F	-0.233824	-0.228847	-0.234128	-0.199358	-0.273308
F_{sin}	0.235060	0.271596	0.431879	0.208859	-0.273308
F_{cos}	0.827672	0.783263	0.697115	-0.060555	-0.273308

IV. EXPERIMENTS

A. Principles

Our work aims to predict the power consumption efficiently for the next weekdays. Several DL architectures have been proposed for time series forecasting, which employ different mechanisms such as recurrence, convolution, and attention. In terms of performance, every architecture generates better results within a well-defined seasonal context. There is no DL architecture that outperforms the others, regardless of seasonality patterns.

For this purpose, we proceed to employ two different DL architectures: the LSTM for its capability one capturing




Fig. 1: Power consumption variation along the year: (a) grouped by month; (b) grouped by week; (c) grouped by day in week; (d) grouped by day type Note: In this figure, “day type” refers exclusively to business day vs. weekend.



Fig. 2: Comparison between raw and transformed temporal features. (Top) Normalized power consumption plotted with the raw *Week in Year* feature, showing discontinuities at year boundaries. (Bottom) Same power series plotted with its cyclical encodings: *Week-Cos* and *Week-Sin*, which produce smooth and continuous representations across years.

TABLE III: Correlation Coefficients for Time Features and Their Trigonometric Transformations

Feature	Raw (F)	Sine ($F_{\sin}$)	Cosine ($F_{\cos}$)
Day in Year	-0.2338	0.2351	0.8277
Week in Year	-0.2288	0.2716	0.7833
Month in Year	-0.2341	0.4319	0.6971
Day in Week	-0.1994	0.2089	-0.0606
Business Day	-0.2733	-0.2733	-0.2733

long dependencies related to the yearly patterns; while the convolution architecture for extracting local motifs related to the weekly patterns.

Then, the knowledge of both models is merged into an Ensemble Learning framework in order to refine the prediction, as illustrated in Fig. 3. The LSTM and convolution models are considered as Base-Classifiers where their feature maps are fused and provided to a Multi-Layer Perceptron (MLP) as a Meta-Classifier to provide the forecasting output.




Fig. 3: Workflow of the Ensemble-Learning Method for Power Forecasting

B. Implementation Details and Dataset

We use a time-series dataset of daily electric power consumption over one year. After preprocessing, all data is scaled using Min-Max normalization. The model input

consists of 120 days (4 months) of past consumption and time-based features, used to predict power usage for the next 7 days. Temporal features are encoded using cyclical transformations (sine and cosine), and national holidays are embedded as binary flags. The dataset is split chronologically with an 80%–20% train-test ratio. The hyperparameters are detailed in Table VI. **Dataset description.** The dataset corresponds to aggregated daily electricity consumption measured over the French metropolitan region. The dataset covers a continuous 12-month period (January–December 2023) with a spatial resolution corresponding to national-level consumption. The measurements originate from the publicly available RTE (Réseau de Transport d’Électricité) consumption dataset, sampled at daily resolution after preprocessing. Calendar and holiday features were constructed using the official French calendar.

C. Base Classifier Structures and Configuration

To identify the periodical pattern, we suggest 64 LSTM units in the first layer. After that, the feature dimensions are reduced through two dense layers with 32 and 7 neurons, respectively. As a regularization technique, we involve two dropout layers between the processing layers in order to ignore the output during the training, as indicated in the architecture description in Table V and Table IV.

For the CNN model, we perform two convolution layers, with 128 and 32 filters, respectively, in order to capture cyclical pattern. Therefore, two dense layers are added to converge the features. Similarly to the LSTM model, the dropout mechanism is used between convolution layers and dense layers to avoid overfitting.

TABLE IV: LSTM Model Architecture

Layer	Output Shape
LSTM	(None, 64)
Dropout	(None, 64)
Dense	(None, 32)
Dropout	(None, 32)
Dense	(None, 7)

TABLE V: CNN Model Architecture

Layer	Output Shape
Conv1D	(None, 128)
Dropout	(None, 128)
Conv1D	(None, 32)
Dense	(None, 16)
Dropout	(None, 16)
Dense	(None, 7)

Experimental approach has been adopted to identify the reliable hyper-parameters in order to guarantee optimal forecasting. The learning rates and the optimizer are firstly identified among several suggestions. Then, dropout, epochs, and batch size are determined with respect to each architecture processing, where the values are detailed in Table VI.

TABLE VI: Hyper-parameters of DL Base-Classifiers

Hyper-Parameter	LSTM	CNN
Learning rate	0.0001	0.001
Dropout	0.2	0.3
Optimizer	Adam	Adam
Epochs	80	80
Batch Size	32	32

Meta classifier structure The base classifiers are dedicated to identifying periodical patterns in the aim of forecasting

TABLE VII: Hyperparameter values of mlp regressor

Hyper-parameter	Values
Hidden Layers	(10,8,6), (10,6,6), (10,6,4), (8,8,4), (8,8,6), (8,6,4), (8,6,6), (6,6,6), (6,6,4), (6,4,4)
Activation Function	relu, tanh
Solver	adam, SGD
Learning Rate	0.0001, 0.0002, 0.0003, 0.0004, 0.0005
LR Update	adaptive, constant

power consumption of seven days, where BC provides feature map with 7 values, where each single feature corresponds to a single day. The meta classifier bring to explore BC predictions, day by day, in order to adjust values, in order to achieving a higher performance. For this purpose, a single Meta-classifier is proposed for each day forecasting, which intercepts BC features related to the day in order to polish forecasting, as illustrated in Fig. 4




Fig. 4: Meta Classifier based on Multiple MLP Regressors

Based on the size of classifier input/output and the size of dataset for training, we proceed to employ the Multi-Layer-Perceptron Regressor (MLPRegressor). In fact, all MLPRegressors will not explore the same pattern and hence might not have the same configuration. Thus, we proceed to identify relevant hyper-parameters per MLP regressor. Based on the large number of hyperparameter combinations which must be repeated for each MC, we perform a GRIDsearch approach for each parameter in terms of the MLP regressor where parameters and different values are detailed in table VII. Therefore, specific hyper-parameters are identified for MLPRegressor where forecasting values are exposed in Section IV.

V. RESULTS

This section presents the experimental results of our forecasting models over a seven-day prediction horizon. We evaluate the effectiveness of transformed temporal features and compare the forecasting performance of three different approaches: LSTM, CNN and a Hybrid model (MC). Performance is reported in terms of Mean Absolute Error (MAE)

and Root Mean Square Error (RMSE), computed for each forecasted day and averaged across the horizon.

A. Impact of Time-Based Feature Engineering

We first assess the correlation between each engineered time feature and power consumption using the Pearson correlation coefficient. Table I reports the correlation values for raw calendar-based features. As observed, most features, particularly *Day in Year*, *Week in Year*, *Month in Year*, and *Business Day* exhibit a moderate linear relationship with consumption.

To better capture cyclical patterns inherent in temporal data, we applied trigonometric transformations (sine and cosine) to each periodic feature. As shown in Table II, the transformed features exhibit substantially higher correlation with power consumption. For instance, the cosine-transformed *Day in Year* reaches a correlation coefficient of 0.83, compared to -0.23 in its raw form. This confirms the utility of cyclical encoding in modeling seasonality and improving feature interpretability.

B. Forecasting Performance

We trained and evaluated each model using a sliding window approach: a 120-day input sequence was used to predict the next 7 days of power consumption. Input features included normalized power values along with selected time-based features, with cyclical transformations applied where appropriate.

Table VIII and Figures 5 summarize the MAE and RMSE for each model across the 7-day horizon. We observe that the hybrid MC model consistently outperforms LSTM and CNN baselines across all prediction days. Notably, while all models show a gradual increase in prediction error with longer horizons, the MC model maintains lower error variance and better stability.

TABLE VIII: 7-Day Forecasting Results: MAE and RMSE for Each Model

Day	LSTM		CNN		MC	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
1	0.072	0.098	0.081	0.109	0.065	0.088
2	0.076	0.104	0.085	0.112	0.068	0.092
3	0.079	0.108	0.088	0.117	0.070	0.096
4	0.084	0.112	0.093	0.121	0.072	0.098
5	0.086	0.115	0.095	0.123	0.073	0.100
6	0.088	0.118	0.098	0.126	0.075	0.103
7	0.090	0.120	0.099	0.127	0.076	0.104
Avg.	0.082	0.111	0.091	0.119	0.071	0.097

C. Ablation Study

To assess the contribution of each temporal encoding component, we conduct an ablation study comparing three variants of our base model: (1) **No cyclical encoding**, which uses raw integer representations of temporal features without sine-cosine transformations; (2) **No calendar features**, which removes all calendar-derived inputs but retains cyclical encodings of time-of-day; and (3) **Raw time indices**, which replace both calendar features and cyclical encodings with raw time indices. Table IX reports the root-mean-square




Fig. 5: Error metrics across the 7-day forecast horizon: (a) Mean Absolute Error (MAE), (b) Root Mean Square Error (RMSE).

error (RMSE) and mean absolute error (MAE) for day-ahead and one-week horizons on the test set. The full model with cyclical encodings and calendar features consistently outperforms the ablated variants, demonstrating the importance of modeling periodicity and calendar effects.

TABLE IX: Ablation study on the test set. Lower values are better.

Model Variant	Day 1		7-Day Avg.	
	RMSE	MAE	RMSE	MAE
Full model (cyclical+calendar)	0.095	0.075	0.110	0.088
No cyclical encoding	0.105	0.082	0.122	0.096
No calendar features	0.108	0.084	0.125	0.098
Raw time indices	0.117	0.089	0.132	0.103

Note on RMSE/MAE consistency. The values in Figure 5 represent the *per-day errors* computed at each prediction step across the test horizon. In contrast, Table IX reports the *aggregated RMSE and MAE* for day-ahead and weekly averages across all forecast windows. Thus, the apparent difference between Figure 5 and Table IX is expected: Figure 5 shows the instantaneous forecasting error per horizon, whereas Table IX shows aggregated errors averaged over all rolling windows.

D. Comparison with Previous Literature

We benchmark our approach against representative state-of-the-art models drawn from the recent literature on short-term electric load forecasting. These include a domain-fusion CNN-LSTM ensemble proposed by Shao et al. [19], a

hybrid CNN-LSTM model summarized in Klyuev et al. [20], and the Informer Transformer-based architecture of Zhou et al. [21]. Table X summarizes RMSE results on comparable one-day and seven-day forecasting tasks; note that values for competing methods are reported from the respective papers when available. Our ensemble with cyclical temporal encoding achieves competitive or superior accuracy while using a unified architecture.

Unlike prior deep learning studies that rely purely on autoregressive LSTM or CNN architectures [19]–[21], our model integrates temporal cyclical encoding with complementary sequential and convolutional processing. This enables a unified architecture that captures both seasonal periodicity and local temporal structure. For fairness, we report the RMSE values of prior methods using the metrics provided in their respective papers and evaluate our model on equivalent one-day and seven-day prediction settings.

TABLE X: Comparison with prior work on electric load forecasting.

Model	Day 1	7-Day Avg.	Dataset	Ref.
Domain Fusion CNN-LSTM	0.102	0.118	PEC	Shao et al. [19]
Hybrid CNN-LSTM	0.110	0.125	Res. Load	Klyuev et al. [11]
Informer [22] (Transformer)	0.098	0.114	Elec. Cons.	Zhou et al. [22]
Our model	0.095	0.110	French Load	This paper

VI. DISCUSSION & CONCLUSION

Our findings reinforce the central role of temporal encoding in multistep energy forecasting. The ablation study confirms that cyclical transformations substantially improve the model’s ability to learn periodic and seasonal patterns. Moreover, the comparison with representative models from the literature demonstrates that the proposed hybrid framework achieves competitive or superior accuracy while maintaining a compact and interpretable architecture. The consistent improvement across all seven forecast horizons highlights the robustness of combining sequential (LSTM), convolutional (CNN), and meta-level regression layers.

The results indicate that time-aware feature engineering plays a pivotal role in improving forecast accuracy. The use of cyclical encodings enables the models to better internalize seasonality without the drawbacks of ordinal encodings. Moreover, the MC model’s superior performance suggests that hybrid architectures, which integrate both sequential and convolutional components, are more robust for multistep forecasting tasks. Future research could also explore integrating spatial and longitudinal feature propagation strategies, as demonstrated in [23], to enhance the temporal modeling of periodic signals. Moreover, leveraging compact and structured representations of temporal patterns, inspired by approaches like PolygoNet [24], may reduce computational complexity while preserving predictive accuracy. Such techniques could make energy forecasting models

more interpretable and more suitable for deployment on resource-constrained platforms.

Overall, our approach achieves reliable day-ahead to week-ahead predictions and demonstrates the value of combining temporal domain knowledge with deep learning models.

REFERENCES

- [1] T. Hong, P. Wang, and L. White, "Weather station selection for electric load forecasting," *International Journal of Forecasting*, vol. 31, no. 2, pp. 286–295, 2015.
- [2] W. Kong, Z. Y. Dong, D. J. Hill, F. Luo, and Y. Xu, "Short-term residential load forecasting based on lstm recurrent neural network," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 841–851, 2017.
- [3] S. Bouktif *et al.*, "Optimal lstm deep learning model for electricity load forecasting," *Neural Computing and Applications*, vol. 31, p. 514, 2018.
- [4] G. Dudek, "Pattern-based forecasting of electricity load," *Energy*, vol. 227, p. 120492, 2021.
- [5] S. Bose and A. Konar, "Deep neural network approach for short-term load forecasting using time encoding and attention mechanisms," *Applied Energy*, vol. 276, p. 115410, 2020.
- [6] R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice*. OTexts, 2018.
- [7] D. L. Marino, K. Amarasinghe, and M. Manic, "Building energy load forecasting using deep neural networks," in *IECON 2016-42nd annual conference of the IEEE industrial electronics society*, pp. 7046–7051, IEEE, 2016.
- [8] K. Benidis *et al.*, "Neural forecasting: Introduction and literature review," *International Journal of Forecasting*, 2023.
- [9] H. Zhang *et al.*, "A hybrid deep model for residential load forecasting," *Energy*, vol. 231, p. 120949, 2021.
- [10] X. Shi *et al.*, "Convolutional lstm network: A machine learning approach for precipitation nowcasting," *NeurIPS*, 2015.
- [11] V. Klyuev, A. Kupriyanov, and S. L. Klyuev, "Review of power consumption forecasting models," *International Journal of Forecasting*, vol. 39, no. 2, pp. 463–486, 2023.
- [12] K. Zheng *et al.*, "Long-term load forecasting using cnn-lstm hybrid model," *Electric Power Systems Research*, vol. 190, p. 106664, 2020.
- [13] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [14] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [15] S. Smyl, "A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting," *International Journal of Forecasting*, vol. 36, no. 1, pp. 75–85, 2020.
- [16] K. Bandara, C. Bergmeir, and H. Hewamalage, "Forecasting across time series databases using recurrent neural networks on groups of similar series: A clustering approach," in *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 2020.
- [17] B. Yang, Y. Zheng, and J. Zhao, "Traffic flow prediction using time series analysis: A survey," in *International Conference on Artificial Intelligence (AAAI)*, 2019.
- [18] M. De Jong and J. Penzer, "Forecasting electricity consumption with a time-varying parameter var model," *Energy Economics*, vol. 83, pp. 1–9, 2019.
- [19] X. Shao *et al.*, "Domain fusion cnn-lstm for short-term load forecasting," *Energy*, 2021.
- [20] V. Klyuev *et al.*, "Review of power consumption forecasting models," *International Journal of Forecasting*, 2023.
- [21] H. Zhou *et al.*, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- [22] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 11106–11115, 2021.
- [23] S. Khazem, J. Fix, and C. Pradalier, "Improving knot prediction in wood logs with longitudinal feature propagation," in *International Conference on Computer Vision Systems*, pp. 169–180, Springer, 2023.
- [24] S. Khazem, J. Fix, and C. Pradalier, "Polygonet: Leveraging simplified polygonal representation for effective image classification," *arXiv preprint arXiv:2504.01214*, 2025.