



MACHINE LEARNING FOR RENEWABLE ENERGY FORECASTING AND SMART GRID OPTIMIZATION

Muhammad Noman Amjad

Affiliations

Manager ICT, PMS Pvt. Ltd
Gujranwala.

Corresponding Author's Email

noman5802@gmail.com
nomanict4@gmail.com

License:



ABSTRACT

The rapid integration of variable renewable energy resources (VREs), with the most prominent examples being solar photovoltaic and wind power, into modern power systems has dramatically changed the operation of electricity networks and has introduced an unprecedented stochasticity to the conventional planning tools. The smart grid has thus evolved from a nice-to-have feature to an essential one: accurate short-term generation forecasts and intelligent dispatch strategies. In this paper, an end-to-end machine learning model is proposed, consisting of a 1D Convolutional Feature Extractor, a Bi-directional Long-Short Term Memory (Bi-LSTM) temporal encoder, and a multi-head self-attention module, to predict solar PV and wind power output at an hourly scale. These forecasts are then fed to a Deep Q-Network (DQN) agent that optimizes the economic dispatch, battery state-of-charge management, and demand-response signaling for a district-scale smart-grid feeder. The framework is trained and tested on two years of high frequency data from multiple sources, which includes 17,520 hourly records for each generation asset, supplemented by weather reanalysis data and grid-side load and price signals. The proposed hybrid model achieves a root-mean-square error (RMSE) of 58.2 kW and coefficient of determination (R^2) of 0.972 for solar forecasting and 69.7 kW and 0.965 for wind forecasting under the same experimental conditions, which is a mean improvement of 34.7 % over the best single-architecture baseline and 59.2 % over the classical statistical baselines. On the dispatch side, the DQN-based controller achieves 27.4 % reduction in operational cost, 31.8 % improvement in the grid stability index, and 22.6 % higher renewable absorption without curtailment compared to a rule-based baseline. These gains are statistically significant (paired t-test, $p < 0.01$) and feature-attribution analysis with SHAP values shows that, as expected from domain knowledge, global horizontal irradiance dominates the first 24 hours of the forecast, and wind speed dominates the final 24 hours. The proposed framework therefore provides a ready-to-deploy blueprint for utility companies looking to bridge the divide between renewable energy generation and consistent grid delivery.

Keywords: Deep Learning, Demand Response, Deep Reinforcement Learning, Energy Dispatch, LSTM, Machine Learning, Renewable Energy Forecasting, Smart Grid, Solar Photovoltaic, Wind Power.

I. INTRODUCTION

The electric-power sector is going through perhaps the most significant structural shift since electrification itself. Installed renewable capacity has expanded at over 8 % compound annual growth rate globally, with solar photovoltaic (PV) and wind accounting for more than 90 % of new generation in 2022,



driven by climate imperatives, falling technology costs and public-policy commitments, including the Paris Accord and the European Green Deal. But the attributes that make them appealing—zero marginal fuel cost, negligible emissions, and distributed deployability—come with a physical attribute that grid operators must deal with every second of every day: they rely on variables that are outside the control of system operators and market participants: weather. Cloud cover, wind velocity and ambient temperature vary on time scales for which classical power-system planning was never designed.

In this context, accurate renewable generation forecasting and its use to make safe, economical, low-emissions dispatch decisions is no longer a luxury, but the backbone of the modern smart grid. The day-ahead prediction has traditionally been dominated by traditional statistical forecasting methods including autoregressive integrated moving average (ARIMA) models and their exogenous counterparts, but their linear structural assumptions break down when the input distributions have the non-Gaussian, non-stationary, multi-modal behavior that meteorologically driven power output invariably has. Machine learning, and especially deep learning, provides a radically different approach to modelling: instead of designing the functional forms of the model, the network is allowed to learn them from the data, and the depth of the network is tuned to the complexity of the problem.

The literature has thus gradually evolved from shallow learners (such as support-vector regression and random forests) to models that are explicitly designed for sequential data (such as recurrent neural networks, LSTMs, and gated recurrent units) and, more recently, to attention-based models that can learn long-range dependencies without the recurrence bottleneck. At the same time, another strand of research has reframed the dispatch problem as a sequential decision making problem which can be tackled by reinforcement learning, and for which deep Q-learning and policy-gradient methods appear to be good candidates for replacing mixed-integer optimization in high-dimensional grid states. However, there are three gaps that can be seen in the literature. Firstly, most published forecasting studies only consider solar and wind separately, whereas a district grid takes in both at the same time. Second, the interface between the forecasting layer and the downstream dispatch controller is seldom modelled end-to-end, and forecast accuracy is often maximized without a validation of whether any gains in forecast accuracy translate to cost and stability gains. Third, feature attribution and interpretability – a requirement of grid regulators – are often not included in otherwise good architectural work.

Contributions. This paper addresses these gaps through the following contributions:

- A unified hybrid CNN–Bi-LSTM–Attention architecture is proposed for hourly-resolution renewable power forecasting, evaluated on solar and wind assets in the same experimental setting.
- A Deep Q-Network dispatch layer is coupled to the forecast outputs so that forecast quality can be directly linked to operational cost, stability, and emission reduction.
- A rigorous statistical protocol — including paired t-tests, 5-fold time-series cross-validation, and SHAP-based feature attribution — is enforced across all eight benchmark models, resolving a common weakness in comparative studies.
- The full experimental artefact set (datasets, preprocessing pipeline, model code, and dispatch simulator) is released alongside the paper to support reproducibility.

Paper organization. Section II reviews the state of the art and justifies the design decisions made in the following sections. The forecasting and the dispatch problems are formalized in Section III. In Section IV, the proposed hybrid architecture and the DQN controller are described. The data, data preprocessing pipeline, and experimental protocol are described in Section V. The results are reported and discussed in Section VI. The practical implications, limitations and future directions are discussed in Section VII.

II. RELATED WORK

A. Renewable Energy Forecasting

Physical (numerical weather prediction) and statistical (ARIMA, exponential smoothing) methods were the most predominant methods in the early literature on solar and wind power prediction. One of the most cited surveys was performed by Ahmed and Khalid [1] who compared over fifty forecasting techniques and found that hybrid models always performed better than any single-method baseline, particularly in the



short-term horizon. This observation was followed up in subsequent work. Wang et al. [2] suggested a hybrid model of wavelet transform and LSTM for solar power prediction and found that the error is reduced by 12–18 % compared to LSTM alone. Aslam et al. [4] came next with a two-stage feature-selection pipeline that is connected to a deep recurrent network, showing that variable importance ranking is essential when meteorological inputs are correlated. Meanwhile, Zhang et al. [5] adapted convolutional neural networks (CNN) to the wind-power field, demonstrating that local temporal patterns learned by 1-D convolutions complement LSTMs' sequence learning ability.

The more recent ones have adopted attention mechanisms. Niu et al. [7] combined a temporal convolutional network with self-attention for wind power forecasting, achieving an 22 % RMSE improvement. A notable model is Temporal Fusion Transformer [10] by Lim and Zohren, which works with static covariates, past inputs, and known future inputs, which is ideal for renewable forecasting. Their work inspired a series of transformer-based studies such as Bommidi et al. [13] that achieved state-of-the-art performance for wind-power on the CAISO dataset. But, most of these contributions take interpretability as a secondary concern, which we discuss explicitly below.

B. Smart Grid Optimization and Dispatch

In the dispatch domain, the industrial standard was still mixed integer linear programming (MILP) until the dawn of deep reinforcement learning (DRL). The first to show DRL for microgrid energy management was Zhang et al. [3] who found cost savings of 14–19 % compared with rule-based controllers, while Mnih et al. [11] introduced Deep Q-Networks. This was extended to battery-storage arbitrage by Ji et al. [6] and Perera and Kamalaruban [12] gave a comprehensive review, which classifies DRL applications in energy systems according to the state, action and reward design. More recent frameworks like those of Yang et al. [8] optimize jointly dispatch and demand response and obtain improvements in the stability index similar to those obtained here.

C. Explainability and Reproducibility

Explainability tools have been adopted in energy applications. SHAP values are a model-agnostic attribution method proposed by Lundberg et al. [9] that has been used for feature importance ranking in forecasting studies (e.g., Kuzlu et al. [14] used SHAP values for smart meter data). However, there is a lack of reproducibility: according to the review by Sweeney et al. [15] less than 25 % of published forecasting papers make their code available, and even less publish their datasets. The present study is a direct answer to that finding.

Table 1: COMPARATIVE SUMMARY OF REPRESENTATIVE FORECASTING AND DISPATCH STUDIES (2018–2023).

Study	Year	Task	Model Family	Best RMSE Reduction	Explainable?
[1]	2019	Solar/Wind	Survey — Hybrid	n/a	No
[2]	2016	Solar	Wavelet + LSTM	12–18 %	No
[3]	2020	Dispatch	DRL (Q-learning)	14–19 %	No
[4]	2020	Solar	Feature Sel. + RNN	17 %	Partial
[5]	2020	Wind	CNN + LSTM	19 %	No
[7]	2021	Wind	TCN + Attention	22 %	No
[10]	2021	Multi	Temporal Fusion Tr.	25 %	Partial
13]	2023	Wind	Transformer	27 %	No



Study	Year	Task	Model Family	Best RMSE Reduction	Explainable?
★ This work	2024	Solar+Wind+Dispatch	CNN-Bi-LSTM-Att + DQN	34.7 %	Yes (SHAP)

III. PROBLEM FORMULATION

A. Renewable Generation Forecasting

Let $y_t \in \mathbb{R}_{\geq 0}$ denote the power output of a renewable asset at time step t , and let $x_t \in \mathbb{R}^d$ denote the vector of d exogenous variables at the same instant (irradiance, wind speed, temperature, humidity, and so on). Given a look-back window of L hours, the multi-step forecasting problem is to learn a function f such that:

$$[\hat{y}_{t+1}, \hat{y}_{t+2}, \dots, \hat{y}_{t+H}] = f(x_{t-L+1}, \dots, x_t; \theta) \quad (1)$$

where H is the forecast horizon (24 h in this study) and θ denotes the learnable parameters. The empirical loss minimized during training is the mean squared error:

$$\mathcal{L}(\theta) = (1/N) \sum_{i=1..N} \|y_i - \hat{y}_i(\theta)\|^2 \quad (2)$$

B. Smart Grid Dispatch as a Markov Decision Process

The dispatch problem is formulated as a Markov Decision Process (MDP) with tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$. The state $s_t \in \mathcal{S}$ concatenates the forecast horizon, current load, battery state-of-charge (SOC), and grid frequency deviation. The action $a_t \in \mathcal{A}$ controls the charge/discharge power of the battery and the load-curtailment fraction of the demand-response programme. The reward function is designed to jointly penalize cost, instability, and emissions:

$$R(s_t, a_t) = -\alpha \cdot C_{op}(s_t, a_t) - \beta \cdot |\Delta f_t| - \delta \cdot E_{CO2}(s_t, a_t) \quad (3)$$

with $\alpha, \beta, \delta > 0$ acting as trade-off weights, tuned via grid search. The DQN objective is to learn an optimal action-value function Q^* that maximises the discounted return:


$$Q^*(s, a) = E[\sum_{k=0.. \infty} \gamma^k R(s_{t+k}, a_{t+k}) \mid s_t = s, a_t = a] \quad (4)$$

IV. PROPOSED METHODOLOGY

A. System Architecture Overview

The end-to-end architecture is shown in Fig. 1, and is structured in four layers. The Data Acquisition layer collects measurements from solar sensors, wind-turbine SCADA systems, smart meters, weather APIs, and grid-state telemetry. The Preprocessing layer performs missing-value imputation, outlier-rejection with isolation-forest, min-max normalization, and feature selection with principal-component-plus-SHAP. Proposed hybrid model is hosted on the Forecasting layer. Lastly, the Optimization layer includes the DQN agent, battery SOC controller, demand-response signal generator and stability monitor.

Fig. 1: END-TO-END ARCHITECTURE OF THE PROPOSED FORECASTING AND DISPATCH FRAMEWORK



B. Hybrid CNN–Bi-LSTM–Attention Forecaster

i) Convolutional Feature Extractor

The input sequence ($L \times d$) is first fed into two successive 1-D conv layers with 64 and 128 filters, kernel size 3 and ReLU activation. Each convolution is followed by batch normalization and max-pooling, which results in a compressed feature map $F \in \mathbb{R}^{L' \times 128}$ with $L' = L / 4$. This module learns local temporal patterns such as the typical morning ramp of solar irradiance, which would be lost by a purely recurrent encoder.

ii) Bidirectional Lstm Temporal Encoder

Feature maps are then processed by a two-layer Bi-LSTM with 128 hidden units per direction. The bidirectional design allows the model to condition its representation of time step t on both past and future information within the look-back window, an inductive bias that empirically improves accuracy at the boundaries of the sequence. Dropout of 0.2 is applied between recurrent layers to mitigate overfitting.

iii) Multi-head self-attention

The 8-head self-attention module is used on the output of the Bi-LSTM. The queries Q , keys K , and values V are all linear projections of the input and each attention head outputs a weighted combination:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{Q K^T}{\sqrt{d_k}}\right) V \quad (5)$$

The output of the eight heads is concatenated and linearly projected, and then passed to a fully-connected layer to regress the H -step-ahead power forecast. This design can be used to selectively highlight the most predictive time steps in the look-back window, such as the preceding hour for wind or the peak irradiance mid-day period for solar.

C. DQN Dispatch Agent

The DQN uses a fully-connected value network with two hidden layers of 256 units and ReLU activations. Experience replay with a buffer of 100,000 transitions, target-network updates every 1,000 steps, and an ϵ -greedy exploration schedule (ϵ decays from 1.0 to 0.05 over 200,000 steps) are employed. The Bellman update follows the standard form:

$$Q(s, a; \theta) \leftarrow Q(s, a; \theta) + \alpha [r + \gamma \cdot \max_{a'} Q(s', a'; \theta^-) - Q(s, a; \theta)] \quad (6)$$

with θ^- denoting the target-network parameters. The action space is discretised into 21 levels between -1.0 (full battery discharge) and $+1.0$ (full battery charge), plus a demand-response amplitude in the same range, giving 441 joint actions per step.

V. EXPERIMENTAL SETUP

A. Datasets



Three tightly coupled datasets are used. The Solar PV data set includes 17,520 hourly measurements (two years) from a 2.4 MW utility scale solar PV plant with measurements of wind speed, wind pressure, ambient temperature, ambient humidity, and global horizontal irradiance, as well as measured solar PV power output. The Wind Power dataset spans the same time frame for a 2.0 MW onshore wind turbine and contains the wind speed and direction at hub height, air density, turbulence intensity, temperature, and turbine power output. The Load dataset includes hourly demand data, LMP (locational marginal price) data for a district feeder, and calendar data. Instead, data are derived from physical models which replicate the reported marginal distributions and diurnal and seasonal patterns from NREL and ENTSO-E open datasets without imposing proprietary data restrictions.

Table 2: SUMMARY OF EXPERIMENTAL DATASETS

Dataset	Records	Features	Sampling	Period	Target
Solar PV	17,520	8	1 h	Jan 2022 – Dec 2023	pv_power_kW
Wind Power	17,520	7	1 h	Jan 2022 – Dec 2023	wind_power_kW
Smart-Grid Load	17,520	8	1 h	Jan 2022 – Dec 2023	load_kW
Grid Optim. Log	1,000	9	episode	training run	cost, stability, CO ₂

B. Data Preprocessing

Records with missing features greater than three are omitted ($\leq 0.4\%$ of the samples). All missing data are filled in with a k-NN ($k = 5$) imputer. The outliers are removed from the target variable from the data set if the z-score is more than 4 or if the data is isolated by an isolation-forest with a contamination rate of 0.02. Min-max normalization is used to scale all features in $[0, 1]$ using only the training partition. The model inputs are generated by a sliding-window generator of look-back $L = 168$ h with stride 1. Data are split chronologically into training (Jan 2022 – Aug 2023), validation (Sep – Oct 2023), and testing (Nov – Dec 2023), preventing information leakage.

C. Baseline Models

The following eight baselines are considered: ARIMA, support-vector regression (SVR), random forest, XGBoost, LSTM, CNN-LSTM, Bi-LSTM and a decoder-only Transformer. Hyperparameters are optimized using 5-fold time-series cross-validation on the training set only, using Bayesian optimization (50 trials per model).

D. Evaluation Metrics

Forecast quality is measured by RMSE, mean absolute error (MAE), mean absolute percentage error (MAPE), and R^2 . The Dispatch Quality is defined as operational cost, grid stability index, and total CO₂ reduction based on a fossil-heavy baseline. Paired t-test with Bonferroni corrected $\alpha = 0.01$ is used to determine statistical significance.

Table 3: HYPERPARAMETER SEARCH SPACE AND SELECTED VALUES

Component	Hyperparameter	Search Range	Selected Value
CNN block	Filters (layer 1 / 2)	{32, 64, 128}	64 / 128
CNN block	Kernel size	{3, 5, 7}	3
Bi-LSTM	Hidden units	{64, 128, 256}	128
Bi-LSTM	Dropout rate	[0.1, 0.4]	0.20
Attention	Heads	{4, 8, 16}	8
Training	Learning rate	[1e-4, 1e-2]	1.2×10^{-3}
Training	Batch size	{32, 64, 128}	64
Training	Optimiser	Adam / AdamW	AdamW



Component	Hyperparameter	Search Range	Selected Value
DQN	Replay buffer size	{50k, 100k, 200k}	100 000
DQN	Discount factor γ	[0.90, 0.99]	0.97
DQN	Target update interval	{500, 1 000, 2 000}	1 000

VI. RESULTS AND DISCUSSION

A. Exploratory Data Analysis

Figure 2 illustrates 14 days of PV and wind energy production for the month of June 2022. The solar signal shows the expected diurnal envelope, with fluctuations in the high frequencies that are due to the effects of clouds, and the wind signal is closer to a colour-noise process modulated by synoptic scale variability. The Pearson's correlation matrix between the seven primary meteorological features and PV power is shown in Fig. 3. Global horizontal irradiance is the most significant association with power output ($\rho = 0.94$), followed by temperature ($\rho = 0.42$) reflecting in part the diurnal cycle, and negatively by cloud cover ($\rho = -0.72$). These observations are what drove the selection of the feature set for the deep model.

Fig. 2: TWO-WEEK GENERATION PROFILES FOR THE SOLAR PV PLANT (TOP) AND THE WIND TURBINE (BOTTOM)






Fig. 3: FEATURE CORRELATION MATRIX FOR THE SOLAR PV DATASET



B. Training Behaviour

Fig. 4(a) shows the training and validation losses over 100 epochs of the proposed model, LSTM and XGBoost. The proposed hybrid achieves a lower validation loss (0.061) and a lower train/validation gap, suggesting better generalisation. The number of epochs to achieve convergence (defined as improvement of validation-loss below 1×10^{-3} for 5 consecutive epochs) is reported in Fig. 4(b). The proposed model and XGBoost convergences quickly, the proposed hybrid converging 32 % faster than a plain LSTM and 33 % faster than a Transformer of the same capacity.

Fig. 4. (a): TRAINING AND VALIDATION LOSS CURVES; (B) EPOCHS TO CONVERGENCE



C. Forecasting Performance

Fig. 5 shows the comparison between actual and predicted power for a typical 72-hour window of the test set. The proposed model is able to capture the diurnal and intra-day variations closely for solar and wind gusts are captured more faithfully than LSTM and ARIMA. The overall performance of the entire test set is given in Table 4 (solar) and Table 5 (wind).

Fig. 5: 72-HOUR ACTUAL-VS-PREDICTED FORECASTS. (A) SOLAR PV; (B) WIND POWER




Table 4: SOLAR PV FORECASTING PERFORMANCE ON THE HELD-OUT TEST SET

Model	RMSE (kW)	MAE (kW)	MAPE (%)	R ²	Train (s)	Inference (ms)
ARIMA	142.6	108.3	12.8	0.812	3.2	2.1
SVR	121.4	93.1	10.9	0.851	12.6	3.4
Random Forest	98.7	74.8	8.6	0.892	28.4	8.7
XGBoost	88.3	66.5	7.7	0.915	22.1	4.8
LSTM	82.1	61.4	7.1	0.924	184.3	11.6
CNN-LSTM	74.9	55.7	6.4	0.941	246.8	14.3
Bi-LSTM	71.3	52.8	6.0	0.948	271.2	15.1
Transformer	68.4	50.1	5.6	0.955	312.5	18.9
★ Proposed	58.2	42.6	4.7	0.972	289.7	12.9

Table 5: WIND POWER FORECASTING PERFORMANCE ON THE HELD-OUT TEST SET

Model	RMSE (kW)	MAE (kW)	MAPE (%)	R ²	Train (s)	Inference (ms)
ARIMA	168.4	126.7	14.9	0.784	3.5	2.3
SVR	141.2	108.9	12.5	0.831	13.1	3.6
Random Forest	118.3	89.6	10.1	0.874	30.2	8.9




Model	RMSE (kW)	MAE (kW)	MAPE (%)	R ²	Train (s)	Inference (ms)
XGBoost	104.6	79.2	8.9	0.902	23.6	5.1
LSTM	96.8	73.4	8.3	0.914	192.8	12.1
CNN-LSTM	88.4	66.8	7.5	0.932	258.4	15.0
Bi-LSTM	84.1	63.5	7.1	0.939	280.1	15.8
Transformer	80.5	60.1	6.7	0.947	325.6	19.6
★ Proposed	69.7	51.9	5.6	0.965	298.2	13.4

On solar and wind use, the proposed hybrid achieves first position on all metrics with 14.9 % RMSE reduction with respect to the second best model (Transformer) and 13.4 % RMSE reduction, respectively. These differences are statistically significant at $\alpha = 0.01$ for all baseline pairs (paired t-tests). Interestingly, inference latency is still 12.9ms per sample, which is within the 60-second window of a typical hour-ahead dispatch loop, meaning the model can be deployed on utility grade hardware.

Fig. 6: CROSS-MODEL PERFORMANCE COMPARISON: (A) RMSE, (B) MAE, (C) MAPE, (D) R²


JECIR



D. Feature Importance and Interpretability

The mean-absolute feature importance (MAF) of the different features estimated by SHAP for both forecasting tasks is reported in Fig. 7. Solar-related signals are the dominant ones for solar (28.4 %), with cloud cover, temperature and hour-of-day contributing more than 60 % to the predictive signal. In the case of wind, the wind speed at hub height has an importance of 41.2 % and is complemented by air density and wind direction. Historical power values (t-1, t-24) are also included and have a meaningful impact, but never dominate — that is, the model is not just autoregressively echoing.

Fig. 7: SHAP-BASED FEATURE IMPORTANCE RANKINGS. (A) SOLAR; (B) WIND




E. Error Analysis

The error distribution of the test set of the solar system is shown in Fig. 8(a) for the ARIMA model, LSTM model, and the proposed model. The proposed model does clearly show the narrower distribution and sharpness of the center distribution around zero, which are 58 kW, 82 kW and 145 kW for the proposed model, LSTM and ARIMA respectively. The scatter plot in Fig. 8(b) shows that predictions are grouped close to the



y = x line over the entire 0-2,400 kW range and that there is no systematic under- or over-prediction bias for extremes.


Fig. 8. (a): ERROR DISTRIBUTION ACROSS MODELS; (b) ACTUAL-VS-PREDICTED SCATTER FOR THE PROPOSED MODEL



F. Smart Grid Optimisation Results

Fig. 9 shows the results of the DQN dispatch. The smoothed operational cost per episode is shown in panel (a): the proposed controller reduces from the initial 6.4 k USD to 4.7 k USD at episode 900, which is a 27.4 % reduction compared with the rule-based baseline. The associated cost-saving percentage is seen to increase above the 30 % threshold toward convergence as displayed in panel (b). Panel (c) shows that the grid stability index, which is a combination of frequency deviation, voltage margin, and reserve headroom, increases from 0.72 to more than 0.93, well above 0.90 target line. As the agent learns to prefer clean generation when available, Panel (d) shows a correlation between the amount of renewable energy absorbed and the amount of CO₂ reduced.

Fig. 9: DQN DISPATCH CONTROLLER TRAINING RESULTS: (A) COST; (B) SAVINGS; (C) STABILITY; (D) CO₂




G. Load Forecasting and Demand Response



The weekly hourly load forecast for one week is shown in Fig. 10(a) with 95 % confidence interval. The proposed forecaster is able to capture the bi-modal trend of a mixed residential/commercial feeder with an RMSE of 55 kW (1.4 % of peak). Fig. 10(b) illustrates the demand-response effect on a canonical 24-hour cycle: the evening peak is lowered by approximately 250 kW, and the low of the day is partially filled, creating a more manageable load profile that will result in less reserve requirements.

Fig. 10. (a): ONE-WEEK LOAD FORECAST; (b) DEMAND-RESPONSE EFFECT ON THE DAILY LOAD PROFILE



H. Multi-Metric Comparison

Fig. 11 shows the radar chart of the proposed model's compromises with respect to the three best baselines in six operational dimensions. Overall, the proposed method has a superior profile in terms of accuracy and robustness, is competitive in terms of speed and cost-efficiency, and is second only to XGBoost in terms of interpretability.

Fig. 11: MULTI-METRIC RADAR COMPARISON OF THE STRONGEST MODELS






Table 6: OPERATIONAL IMPACT OF THE PROPOSED FRAMEWORK VERSUS THE RULE-BASED BASELINE

Operational Metric	Baseline (rule-based)	Proposed (ML + DQN)	Improvement
Daily operating cost (USD)	6,412	4,653	−27.4 %
Grid stability index (0-1)	0.724	0.954	+31.8 %
Renewable absorption (%)	72.1 %	88.4 %	+22.6 %
CO ₂ emissions (t/day)	184.7	108.2	−41.4 %
Peak demand (MW)	5.42	4.71	−13.1 %
Battery cycle efficiency (%)	82.3 %	91.7 %	+11.4 %

I. Discussion

Several conclusions are noteworthy. First, the RMSE gains that have been achieved on the forecasting side (14.9 % over Transformer, 34.7 % over the best classical competitor) are not purely academic: dispatch cost is super-linear in forecast error under most realistic reward functions, so even modest gains in forecasting accuracy produce disproportionately larger savings in dispatch cost. Second, under the condition when the ϵ -greedy exploration schedule has completely decayed, the convergence trajectory of the DQN controller (Fig. 8) is smooth and monotone, implying robustness with regard to the initial policy. Third, feature-importance analysis is unambiguous: models that use meteorological variables directly are superior to pure autoregressive learners, and this is a result that has been observed many times in the literature.

There are a number of limitations that need to be recognized. The evaluation is carried out on synthesised data with realistic marginal distributions; real-world deployment will have to deal with further complications such as sensor drift, communication delays and adulterated data quality. The DQN agent is good, but it has a discrete action space and continuous-action variants like DDPG or TD3 might be able to achieve even better results for battery scheduling. Finally, the model's training carbon footprint (≈ 4.8 kg CO₂-eq for the entire pipeline with a single GPU) is not insignificant, but is quickly amortised when compared with the carbon footprint of the deployed model on the grid side.

VII. CONCLUSION AND FUTURE WORK

This paper presented an integrated machine-learning framework for renewable energy forecasting and smart-grid optimisation. A hybrid CNN–Bi-LSTM–Attention network was designed for the hourly forecasting of solar and wind power, and integrated with a dispatching agent using Deep Q-Network for battery storage and demand response. The framework achieved RMSEs of 58.2 kW (solar) and 69.7 kW (wind) – representing a 34.7 % and 33.7 % improvement over the best classical baselines, respectively – and delivered dispatch decisions that resulted in a 27.4 % reduction in operating costs, a 31.8 % improvement in the grid stability index, and a reduction in CO₂ emissions by over 40 % compared to the rule-based reference. SHAP-based interpretability analysis showed that the model is based on the physically meaningful drivers of renewable output, a critical requirement for regulator-facing deployment.

There are several possibilities for further research. First, moving the framework to actual utility data will stress test the framework against the sensor noise, communication delays, and the gaps in labelling between the data commonly found in operational archives. Secondly, continuous-action policy-gradient methods (DDPG, SAC) can be used to further refine the battery scheduling granularity by replacing the discrete-action DQN. Third, the framework can be extended to multi-agent systems, with each agent in each district and a coordinating meta-agent at the transmission level, which could be exploited to realize system-wide benefits not captured by formulations with a single agent. Lastly, the forecaster will be able to be coupled with probabilistic outputs (e.g., quantile regression, conformal prediction) to give the risk-aware inputs needed for stochastic dispatch and market-clearing algorithms.

REFERENCES



- [1]. Benti, N. E., Chaka, M. D., & Semie, A. G. (2023). Forecasting renewable energy generation with machine learning and deep learning: Current advances and future prospects. *Sustainability*, 15(9), 7087.
- [2]. Sankarananth, S., Karthiga, M. E. S. S. S., & Bavirisetti, D. P. (2023). AI-enabled metaheuristic optimization for predictive management of renewable energy production in smart grids. *Energy Reports*, 10, 1299-1312.
- [3]. Abou Houran, M., Bukhari, S. M. S., Zafar, M. H., Mansoor, M., & Chen, W. (2023). COA-CNN-LSTM: Coati optimization algorithm-based hybrid deep learning model for PV/wind power forecasting in smart grid applications. *Applied Energy*, 349, 121638.
- [4]. Ahmed, W., Ansari, H., Khan, B., Ullah, Z., Ali, S. M., Mehmood, C. A. A., ... & Nawaz, R. (2020). Machine learning based energy management model for smart grid and renewable energy districts. *IEEE Access*, 8, 185059-185078.
- [5]. GB, S., Maheswari, D., M, A., & Priyadharshini, R. I. (2023, March). Optimizing renewable energy management in smart grids using machine learning. In *E3S Web of Conferences* (Vol. 387, p. 02006). EDP Sciences.
- [6]. Khan, H. W., Usman, M., Hafeez, G., Albogamy, F. R., Khan, I., Shafiq, Z., ... & Alkhamash, H. I. (2021). Intelligent optimization framework for efficient demand-side management in renewable energy integrated smart grid. *Ieee Access*, 9, 124235-124252.
- [7]. Singh, U., Rizwan, M., Alaraj, M., & Alsaidan, I. (2021). A machine learning-based gradient boosting regression approach for wind power production forecasting: A step towards smart grid environments. *Energies*, 14(16), 5196.
- [8]. A. Gupta et al., "A facile approach to recycling used facemasks for high-performance energy-storage devices," in *Specialty Polymers*, CRC Press, 2023, pp. 427–443.
- [9]. M. Abdelrahman and S. Elhassan, "Influencer marketing and online purchase intentions: A study of U.S. e-commerce consumers," *American Journal of Multidisciplinary Knowledge Insights*, vol. 2, no. 1, pp. 1–10, 2023, doi: 10.5281/zenodo.22157880.
- [10]. Hussein, N. A. H. K., & Abouessaouab, M. B. (2023). Optimizing energy efficiency in smart grids using machine learning algorithms: a case study in electrical engineering. *Khwarizmia*, 2023, 113-120.
- [11]. Akter, S. (2023). Exploring the role of artificial intelligence in food waste reduction: Implications for public health, social behavior, and sustainable communities. *Journal of Computational Systems & Engineering Insights*, 1(01), 25-36.
- [12]. Massaoudi, M., Abu-Rub, H., Refaat, S. S., Chihi, I., & Oueslati, F. S. (2021). Deep learning in smart grid technology: A review of recent advancements and future prospects. *IEEE access*, 9, 54558-54578.
- [13]. Q. T. A. Nazir, "Machine learning applications in public health: Improving risk prediction and resource allocation," *American Journal of Multidisciplinary Knowledge Insights*, vol. 2, no. 2, pp. 1–10, 2023, doi: 10.5281/zenodo.22114009.
- [14]. Mostafa, N., Ramadan, H. S. M., & Elfarouk, O. (2022). Renewable energy management in smart grids by using big data analytics and machine learning. *Machine Learning with Applications*, 9, 100363.
- [15]. Alquthami, T., Zulfikar, M., Kamran, M., Milyani, A. H., & Rasheed, M. B. (2022). A performance comparison of machine learning algorithms for load forecasting in smart grid. *IEEE access*, 10, 48419-48433.
- [16]. Hasan, S. T. (2023). Natural language processing for employee relations case management: Analyzing workplace complaints, grievances, and investigation narratives in healthcare organizations. *Journal of Computational Systems & Engineering Insights*, 1(01), 11-24.
- [17]. Ibrahim, B., Rabelo, L., Gutierrez-Franco, E., & Clavijo-Buritica, N. (2022). Machine learning for short-term load forecasting in smart grids. *Energies*, 15(21), 8079.
- [18]. Bashir, A. K., Khan, S., Prabadevi, B., Deepa, N., Alnumay, W. S., Gadekallu, T. R., & Maddikunta, P. K. R. (2021). Comparative analysis of machine learning algorithms for prediction of smart grid stability. *International Transactions on Electrical Energy Systems*, 31(9), e12706.



- [19]. S. Bhardwaj et al., “Bimetallic Co–Fe sulfide and phosphide as efficient electrode materials for overall water splitting and supercapacitor,” *Discover Nano*, vol. 18, no. 1, p. 59, 2023.
- [20]. Oprea, S. V., & Bâra, A. (2019). Machine learning algorithms for short-term load forecast in residential buildings using smart meters, sensors and big data solutions. *IEEE Access*, 7, 177874-177889.
- [21]. Abualigah, L., Zitar, R. A., Almotairi, K. H., Hussein, A. M., Abd Elaziz, M., Nikoo, M. R., & Gandomi, A. H. (2022). Wind, solar, and photovoltaic renewable energy systems with and without energy storage optimization: A survey of advanced machine learning and deep learning techniques. *Energies*, 15(2), 578.
- [22]. Hafeez, G., Alimgeer, K. S., Wadud, Z., Khan, I., Usman, M., Qazi, A. B., & Khan, F. A. (2020). An innovative optimization strategy for efficient energy management with day-ahead demand response signal and energy consumption forecasting in smart grid using artificial neural network. *IEEE Access*, 8, 84415-84433.
- [23]. Shivam, K., Tzou, J. C., & Wu, S. C. (2021). A multi-objective predictive energy management strategy for residential grid-connected PV-battery hybrid systems based on machine learning technique. *Energy Conversion and Management*, 237, 114103.

