RESEARCH ARTICLE

A Hybrid CNN-LSTM Deep Learning Framework for Demand Response Management in Grid-Connected Renewable Energy Microgrids

Md. Ahsan Habib[1], **Md. Sadik Hassan Arik**[1], **Md. Sabiruzzaman**[1] and **Md. Nimul Hasan**[2,3,*]

[1] Department of Industrial Engineering, Lamar University, Beaumont TX 77710, United States
[2] Department of Electrical, Electronic and Communication Engineering, Pabna University of Science and Technology, Pabna 6600, Bangladesh
[3] School of Science and Engineering, Uttara University, Dhaka 1230, Bangladesh

Abstract

Grid-connected microgrids integrating solar photovoltaic and wind generation face a persistent mismatch between intermittent supply and local demand, raising operating cost and increasing reliance on grid imports during peak periods. Demand response (DR) mitigates this mismatch by reshaping load, but its effectiveness depends on accurate forecasting of load and renewable generation. This paper proposes a hybrid CNN-LSTM deep learning framework for DR management in a grid-connected microgrid comprising solar photovoltaic, wind, and battery storage. Convolutional layers extract local temporal features, long short-term memory layers model long-range daily dependencies, and the resulting forecasts drive a mixed-integer DR optimization layer that schedules flexible loads and battery dispatch to minimize operating cost. The framework is evaluated on one year of hourly HOMER Pro data for a residential community in Mirpur, Dhaka (2024), with the DR layer assessed on 24-hour scheduling days and benchmarked against standalone CNN, LSTM, and GRU models, as well as no-DR and rule-based DR strategies. The proposed model lowers load mean absolute percentage error to 4.8%, versus 7.0% for the next-best baseline. The DR layer reduces daily operating cost by 28.8% relative to no-DR and by 15.9% relative to rule-based DR, lowers the peak-to-average ratio from 1.62 to 1.24, and raises renewable utilization to 76.8%. Sensitivity analyses indicate gains across flexible load share, battery capacity, and forecast horizon, and illustrative seasonal days show similar behavior. Isolating the share of gains due to forecast accuracy remains future work.

Keywords: demand response, deep learning, CNN-LSTM, microgrid, renewable energy, load forecasting.

Submitted: 16 June 2026
Revised: 25 October 2026
Accepted: 26 October 2026
Published: 28 November 2026

Vol. 1, **No.** 1, 2026.
10.62762/JISES.2026.848280

***Corresponding author:**
✉ Md. Nimul Hasan
nimul34.din.h@gmail.com

Citation
Habib, M. A., Arik, M. S. H., Sabiruzzaman, M., & Hasan, M. N. (2026). A Hybrid CNN-LSTM Deep Learning Framework for Demand Response Management in Grid-Connected Renewable Energy Microgrids. *Journal of Intelligent and Sustainable Energy Systems*, 1(1), 18–32.

© 2026 by the Authors. Published by Institute of Central Computation and Knowledge. This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/).

1 Introduction

Grid-connected microgrids integrate distributed renewable resources such as solar photovoltaic and wind generation with conventional grid supply to serve local loads. The intermittent nature of these resources creates a mismatch between generation and demand, increasing reliance on grid imports during peak periods and raising operational costs. Demand response (DR) addresses this mismatch by shifting, curtailing, or reshaping load consumption patterns in response to grid signals, generation availability, or pricing schemes. Effective DR coordinates flexible loads and battery storage to reduce peak demand and operating cost [1].

Optimal dispatch of renewable-based urban microgrids depends on accurate forecasts of electrical load and wind power, which deep learning models can supply [2]. Accurate, real-time forecasting of load and generation is essential for DR programs to function effectively. Deep learning models capture temporal patterns in energy data and are widely reported to outperform statistical and shallow machine learning approaches in load and renewable generation forecasting tasks [3].

Convolutional Neural Networks (CNNs) extract local temporal features from input data, while Long Short-Term Memory (LSTM) networks capture long-range temporal dependencies in sequential data. A hybrid CNN-LSTM architecture combines these strengths: the CNN layer extracts short-term feature patterns from raw time-series data, and the LSTM layer models the temporal evolution of these features. This hybrid structure improves short-term load forecasting accuracy [4, 5], which directly supports more effective DR decision-making.

This paper proposes a hybrid CNN-LSTM deep learning framework for demand response management in a grid-connected microgrid comprising solar PV, wind generation, battery storage, and grid interconnection. The framework forecasts load demand and renewable generation, then uses these forecasts to drive a DR optimization layer that schedules controllable loads and battery dispatch to minimize grid dependency and operational cost.

The main contributions of this paper are:

- A hybrid CNN-LSTM model for short-term load and renewable generation forecasting in a grid-connected microgrid.
- A DR optimization strategy that uses forecast outputs to schedule flexible loads and battery operation.
- A simulation-based evaluation comparing the proposed framework against baseline forecasting and DR methods across multiple performance metrics.

The remainder of this paper is organized as follows. Section 2 reviews related work on demand response, deep learning-based forecasting, and hybrid model architectures. Section 3 describes the microgrid system architecture. Section 4 presents the proposed hybrid CNN-LSTM model and DR optimization formulation. Section 5 describes the simulation setup. Section 6 presents results and discussion. Section 7 concludes the paper.

2 Literature Review

2.1 Demand Response in Renewable-Integrated Microgrids

Demand response programs in microgrids fall into two categories: price-based programs, where consumers adjust usage based on time-varying tariffs, and incentive-based programs, where utilities directly control or compensate for load adjustments. A mixed-integer linear programming (MILP) framework for DR in solar-powered microgrids jointly optimizes load classification, dynamic price thresholding, and battery storage utilization across multiple periods [1]. Coordinated wind, storage, and DR strategies reduce operating cost and improve peak load management in wind-integrated microgrids [6]. Deep reinforcement learning approaches that prioritize flexible demand-side resources, further improve microgrid dispatch under uncertainty [7]. However, these DR frameworks take load and renewable generation profiles as inputs, so the quality of their schedules depends on the accuracy of those profiles under weather-driven variability.

2.2 Deep Learning for Load and Renewable Generation Forecasting

Recurrent architectures, particularly LSTM and Gated Recurrent Unit (GRU) networks, are widely applied to short-term load forecasting due to their ability to model sequential dependencies in time-series data [3]. A hybrid GRU and Temporal Convolutional Network (TCN) model with an attention mechanism captures both long-term dependencies and local patterns in load data, improving prediction performance over standalone architectures [8]. For renewable generation

forecasting, deep learning models that combine load and wind power prediction support optimal dispatch of urban microgrids with PV, wind, and battery storage [2]. Machine learning algorithms applied to historical production data and weather patterns achieve low forecasting error for solar and wind power in grid-connected microgrids [9].

2.3 Hybrid CNN-LSTM Architectures

Hybrid CNN-LSTM models combine convolutional feature extraction with LSTM-based temporal modeling. A parallel multichannel network using one-dimensional CNN and bidirectional LSTM layers improves short-term load forecasting accuracy in smart grids [4]. A multi-task CNN-LSTM architecture that shares parameters across forecasting tasks reduces mean squared error, root mean squared error, and mean absolute error compared to single-task models [5]. A parallel CNN-LSTM model with a soft attention mechanism, evaluated on AEP and ISO New England datasets, achieves lower RMSE, MAE, and MAPE than the published models it is compared with, in both single-step and multi-step forecasting [10]. A one-dimensional CNN-LSTM model enhanced with integrated gradients attribution and explainable AI (IGA-XAI) improves day-ahead peak and off-peak load forecasting and generalizes across multiple cities [11].

2.4 Integration of Forecasting with Demand Response Optimization

Recent work links forecasting outputs directly to DR optimization layers, where forecasted load and generation values become inputs to scheduling algorithms for controllable loads, energy storage, and grid exchange. Deep learning-based demand response models for short-term microgrid operation simultaneously minimize operational cost and emissions in the presence of distributed energy resources [12]. On the optimization side, a price-elastic DR approach combined with a greedy rat swarm optimization algorithm improves both economic and environmental performance of microgrids under real-time pricing (RTP), time-of-use (TOU), and critical peak pricing (CPP) schemes [13]. Because such optimization layers take load and generation profiles as inputs, forecasting accuracy is expected to affect the quality of DR scheduling decisions, which motivates the use of high-accuracy hybrid deep learning models as the forecasting backbone.

2.5 Summary of Reviewed Studies

Table 1 summarizes the methods, application domains, and key findings of the studies discussed above.

Table 1. Summary of reviewed studies on deep learning forecasting and demand response in microgrids.

Ref.	Method	Application	Key Finding
[1]	MILP	DR scheduling, solar microgrid with battery	Joint optimization of load classification, price thresholding, and battery state-of-charge constraints
[2]	Hybrid deep learning + MILP	Load and wind forecasting, urban microgrid dispatch	Hybrid model improves 3-day-ahead load and wind power forecasts over statistical baselines
[3]	Survey	Load and renewable forecasting, smart microgrids	Identifies deep learning as the dominant approach for power load and renewable forecasting
[4]	1D CNN + Bi-LSTM (parallel)	Short-term load forecasting	Parallel multichannel structure improves forecasting accuracy over single-branch models
[5]	Multi-task CNN-LSTM	Short and medium-term load forecasting	Multi-task learning reduces MSE, RMSE, and MAE versus single-task CNN-LSTM
[8]	GRU + TCN + attention	Short-term load forecasting, smart grid	Attention mechanism focuses on relevant input segments, improving prediction accuracy
[9]	Machine learning-based forecasting	Power generation forecasting, grid-connected microgrid	Low MSE for solar and wind forecasts; supports DR and energy management
[10]	Parallel CNN-LSTM with soft attention	Short-term load forecasting, smart grid	Lowest RMSE, MAE, and MAPE on AEP and ISO New England data versus the published models compared
[11]	1D CNN-LSTM with IGA-XAI	Day-ahead to 7-day load forecasting	R^2 up to 0.96 across peak, evening peak, and off-peak periods; generalizes across cities
[12]	Deep learning-based DR model	Short-term renewable microgrid operation	Minimizes operational cost and emissions with distributed energy resources
[13]	Price-elastic DR + greedy rat swarm optimization	Economic and environmental microgrid optimization	Improves cost and emission outcomes under RTP, TOU, and CPP pricing schemes
[7]	Deep reinforcement learning	Priority-based flexible demand-side dispatch	Coordinates wind, storage, and flexible loads to address demand-side uncertainty
[6]	Coordinated storage + hybrid DR	Day-ahead economic dispatch, wind-integrated microgrid	Combined wind, storage, and DR significantly reduces operating cost and improves peak management

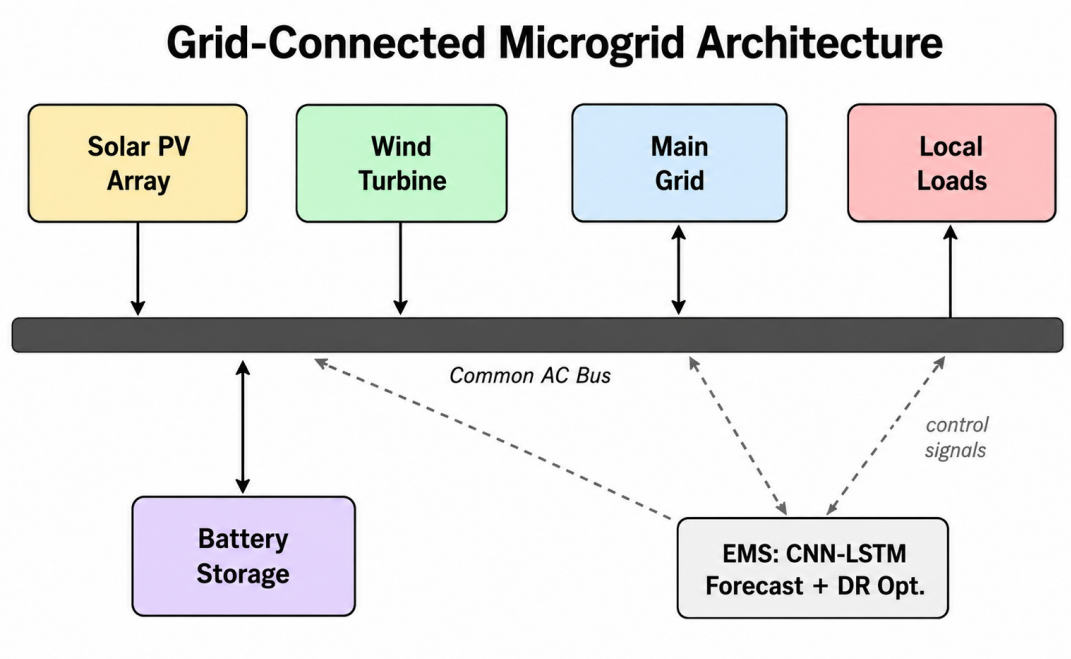

Figure 1. Architecture of the grid-connected microgrid with solar PV, wind generation, battery storage, local loads, and the energy management system (EMS).

2.6 Research Gap

Many existing studies address forecasting and DR separately, or apply standalone deep learning models without exploiting hybrid architectures suited to both local and long-range temporal patterns in microgrid data. Limited work integrates a CNN-LSTM forecasting model directly into a DR optimization framework for grid-connected microgrids with solar, wind, and battery storage. This paper addresses this gap by proposing an integrated hybrid CNN-LSTM forecasting and DR optimization framework, evaluated through simulation on a grid-connected microgrid model.

3 Microgrid System Architecture

3.1 System Overview

The studied system is a grid-connected microgrid comprising four power components: a solar photovoltaic (PV) array, a wind turbine generator, a battery energy storage system (BESS), and a bidirectional connection to the main utility grid. These components supply a set of local loads through a common AC bus. The main grid acts as a flexible source and sink, importing power during generation deficits and exporting surplus generation when local supply exceeds demand [14]. The battery provides short-term flexibility by storing surplus renewable generation and discharging during peak or high-tariff periods [15]. Figure 1 shows the overall architecture.

3.2 Generation Components

The PV array output depends on solar irradiance and ambient temperature. The instantaneous PV power is modeled as

$$P_{pv}(t) = \eta_{pv}\, A_{pv}\, G(t)\, \left[1 - \beta\,(T_c(t) - T_{ref})\right], \qquad (1)$$

where η_{pv} is the panel efficiency, A_{pv} is the array area, $G(t)$ is the solar irradiance, β is the temperature coefficient, $T_c(t)$ is the cell temperature, and T_{ref} is the reference temperature.

The wind turbine output is a function of wind speed $v(t)$ and follows the standard piecewise power curve

$$P_w(t) = \begin{cases} 0, & v(t) < v_{ci} \text{ or } v(t) \geq v_{co} \\ P_r \dfrac{v(t)^3 - v_{ci}^3}{v_r^3 - v_{ci}^3}, & v_{ci} \leq v(t) < v_r \\ P_r, & v_r \leq v(t) < v_{co} \end{cases} \qquad (2)$$

where v_{ci}, v_r, and v_{co} are the cut-in, rated, and cut-out wind speeds, and P_r is the rated turbine power.

3.3 Battery Energy Storage System

The battery state of charge (SOC) evolves according to the charge and discharge power at each time step

$$SOC(t+1) = SOC(t) + \frac{\eta_{ch}\, P_{ch}(t) - P_{dis}(t)/\eta_{dis}}{E_{cap}}\, \Delta t, \qquad (3)$$

where $P_{ch}(t)$ and $P_{dis}(t)$ are the charge and discharge power, η_{ch} and η_{dis} are the charge and discharge

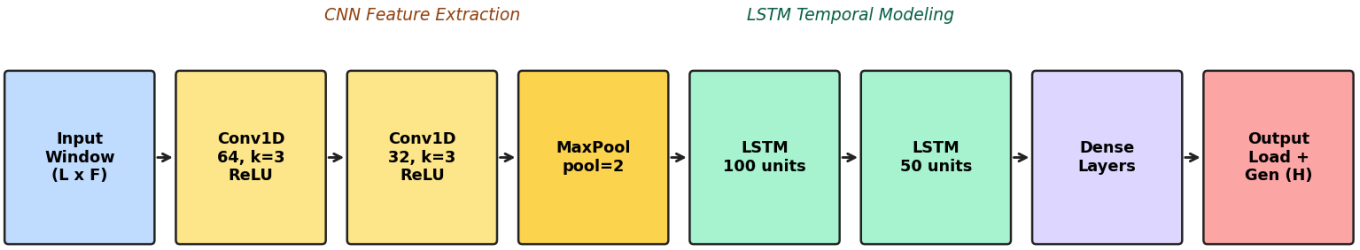

Figure 2. Structure of the proposed hybrid CNN-LSTM forecasting model with the tensor size at each layer. F is the number of input features.

efficiencies, E_{cap} is the battery capacity, and Δt is the time step. The SOC is bounded by $SOC_{min} \leq SOC(t) \leq SOC_{max}$ to limit degradation, and charge and discharge cannot occur simultaneously [15].

3.4 Power Balance and Grid Exchange

At each time step, total supply must equal total demand. The power balance constraint is

$$P_{pv}(t) + P_{w}(t) + P_{dis}(t) + P_{grid}^{imp}(t) = P_{load}(t) + P_{ch}(t) + P_{grid}^{exp}(t), \quad (4)$$

where $P_{load}(t)$ is the total load demand, and $P_{grid}^{imp}(t)$ and $P_{grid}^{exp}(t)$ are the power imported from and exported to the main grid. Grid exchange is bounded by the rated capacity of the point of common coupling [14].

3.5 Energy Management and Demand Response Layer

A two-layer control structure governs operation. The upper layer runs the hybrid CNN-LSTM forecasting model, which predicts load demand and renewable generation over the scheduling horizon. The lower layer takes these forecasts and solves the demand response optimization problem, dispatching the battery, scheduling flexible loads, and setting grid exchange set-points to minimize operating cost [14]. Loads are classified as fixed or flexible; flexible loads can be shifted within comfort and operational constraints to reduce peak demand and improve renewable utilization. Section 4 details the forecasting model and the optimization formulation.

4 Proposed Methodology

The proposed framework consists of two stages: a hybrid CNN-LSTM forecasting model that predicts load demand and renewable generation, and a demand response optimization layer that uses these forecasts to schedule flexible loads and battery dispatch. Figure 2 shows the structure of the forecasting model.

4.1 Input Representation and Preprocessing

The model input is a sliding window of historical multivariate time-series data. At time t, the input window of length L is

$$\mathbf{X}_t = [\mathbf{x}_{t-L+1}, \mathbf{x}_{t-L+2}, \ldots, \mathbf{x}_t], \quad (5)$$

where each $\mathbf{x}_\tau$ contains load demand, solar irradiance, wind speed, ambient temperature, and time-of-day and day-of-week indicators. The day-of-week indicators carry little information in this dataset, because the load derives from a single 24-hour record (Section 5.2). We keep them so that the model applies to measured data with weekly patterns. All features are normalized to the range $[0, 1]$ using min-max scaling to stabilize training. The model predicts the load and renewable generation for the next H steps over the scheduling horizon.

4.2 Convolutional Feature Extraction

The first stage applies one-dimensional convolution to the input window to extract local temporal features. For a convolutional layer with filter $\mathbf{w}_k$ and bias b_k, the feature map is

$$c_k(i) = \mathrm{ReLU}\left(\sum_{j=0}^{K-1} w_k(j)\, x(i+j) + b_k\right), \quad (6)$$

where K is the kernel size and $\mathrm{ReLU}(z) = \max(0, z)$ is the activation function. The proposed model stacks two Conv1D layers with 64 and 32 filters and kernel size 3, each capturing short-term trends and cyclical patterns in the data [10, 16]. A max-pooling layer with pool size 2 follows the convolutional layers to reduce dimensionality and retain the most salient features.

4.3 LSTM Temporal Modeling

The pooled feature maps form a sequence that is passed to the LSTM layers, which model long-range

temporal dependencies. Each LSTM cell updates its state through the forget, input, and output gates

$$f_t = \sigma\left(W_f[h_{t-1}, x_t] + b_f\right), \tag{7}$$
$$i_t = \sigma\left(W_i[h_{t-1}, x_t] + b_i\right), \tag{8}$$
$$\tilde{C}_t = \tanh\left(W_C[h_{t-1}, x_t] + b_C\right), \tag{9}$$
$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t, \tag{10}$$
$$o_t = \sigma\left(W_o[h_{t-1}, x_t] + b_o\right), \tag{11}$$
$$h_t = o_t \odot \tanh(C_t), \tag{12}$$

where σ is the sigmoid function, $\odot$ denotes element-wise multiplication, C_t is the cell state, and h_t is the hidden state. The proposed model uses two stacked LSTM layers with 100 and 50 units. A dropout of 0.2 is applied after each LSTM layer to reduce overfitting.

4.4 Output Layer and Training

The final LSTM hidden state passes through two fully connected dense layers: a 64-unit hidden layer and a linear output layer of 48 units. The output is a single joint vector, reshaped to 24×2 values, that holds the 24 hourly load forecasts and the 24 hourly renewable-generation forecasts over the scheduling horizon. The model is trained by minimizing the mean squared error between predicted and actual values,

$$\mathcal{L} = \frac{1}{N}\sum_{n=1}^{N}\left(y_n - \hat{y}_n\right)^2, \tag{13}$$

where y_n and $\hat{y}_n$ are the actual and predicted values. The Adam optimizer [17] with a learning rate of 0.001 is used, and training runs for up to 100 epochs with early stopping on validation loss.

The architecture therefore performs direct multi-output forecasting rather than recursive one-step-ahead prediction. No predicted value is fed back as an input, which reduces the accumulation of forecasting errors across later steps of the horizon. Because the two Conv1D layers use padding, the 24-step input keeps its length through the convolutions and becomes a sequence of 12 feature vectors of dimension 32 after max-pooling. The horizon-error curve in Figure 8 reports the error at each step h of this output, averaged over all test windows.

4.5 Demand Response Optimization

The forecasted load and generation values become inputs to the DR optimization layer, which schedules flexible loads, battery dispatch, and grid exchange to minimize total operating cost over the horizon H. The objective function is

$$\min \sum_{t=1}^{H}\left[\lambda_t^{imp} P_{grid}^{imp}(t) - \lambda_t^{exp} P_{grid}^{exp}(t) + \lambda^{deg} P_{dis}(t)\right]\Delta t, \tag{14}$$

where λ_t^{imp} and λ_t^{exp} are the time-varying grid import and export prices, and λ^{deg} is the battery degradation cost coefficient. The optimization is subject to the power balance constraint in Eq. (4), the SOC dynamics in Eq. (3), and the battery, grid, and flexible load limits. This linear throughput model ignores the effect of depth of discharge, C-rate, and temperature on aging. With a 200 kW power limit on an 800 kWh battery, the maximum C-rate is 0.25C, a mild operating range. A cycle-counting degradation model is left for future work.

Flexible loads are scheduled within a permitted shifting window. Each shiftable load l has rated power P_l, duration d_l hours, and allowed window $[a_l, b_l]$. A binary variable $s_l(t)$ marks the hour in which the load

Table 2. Microgrid simulation parameters.

Parameter	Symbol	Value
PV rated capacity	P_{pv}^{rated}	500 kW
Wind rated capacity	P_r	400 kW
Battery capacity	E_{cap}	800 kWh
Battery power limit	P_{batt}^{max}	200 kW
SOC limits	$[SOC_{min}, SOC_{max}]$	[0.2, 0.9]
Charge/discharge efficiency	η_{ch}, η_{dis}	0.95
Grid exchange limit	P_{grid}^{max}	600 kW
Design peak load	P_{load}^{peak}	700 kW
Wind cut-in / rated / cut-out	v_{ci}, v_r, v_{co}	3 / 12 / 25 m/s
Time-of-use peak price	λ^{peak}	0.18 $/kWh
Time-of-use off-peak price	λ^{off}	0.06 $/kWh

starts:

$$\sum_{t=a_l}^{b_l-d_l+1} s_l(t) = 1, \qquad s_l(t) \in \{0,1\} \tag{15}$$

$$u_l(t) = \sum_{\tau=t-d_l+1}^{t} s_l(\tau) \tag{16}$$

$$P_{load}(t) = P_{fixed}(t) + \sum_{l} P_l\, u_l(t) \tag{17}$$

where $P_{fixed}(t)$ is the non-shiftable load and $s_l(\tau) = 0$ outside the allowed start hours. The load runs for d_l contiguous hours, and its energy equals $D_l = P_l d_l$. The formulation remains linear, so the problem stays a mixed-integer linear program. The results in Section 6 were obtained with the single-slot form, in which the whole demand D_l sits in one hour ($d_l = 1$ h). A single-slot load can sit in the cheapest hour, while a multi-hour load must also occupy neighboring hours. The reported savings are therefore likely optimistic for long-duration appliances. This formulation shifts flexible demand from high-price peak periods to periods of high renewable availability or low grid price, reducing operating cost and improving renewable utilization [1, 13].

5 Simulation Setup

5.1 Microgrid Configuration

The microgrid is simulated over a 24-hour horizon at one-hour resolution. The system parameters are listed in Table 2. The PV array and wind turbine are sized to meet a significant share of the daily load, with the battery and grid connection providing flexibility. Flexible loads account for 50% of the total load in the base case. Time-varying grid prices follow a time-of-use tariff with peak, mid-peak, and off-peak periods.

5.2 Dataset and Training

The forecasting model uses hourly data generated with HOMER Pro (Hybrid Optimization of Multiple Energy Resources) for a residential community in Mirpur, Dhaka, Bangladesh, over the year 2024, 1 January to 31 December, a total of 8,760 hourly records. The load derives from a 24-hour residential load record, which HOMER extends over the year with random day-to-day and hour-to-hour variability. Solar irradiance, wind speed, and ambient temperature come from NASA resource data supplied through HOMER. PV and wind power are the HOMER component outputs for the capacities in Table 2. A sliding window with $L = 24$ and $H = 24$ yields 8,713 input-output samples. A chronological split assigns 6,132, 1,314, and 1,314 hourly records (70%, 15%, and 15%) to training, validation, and testing, so the test period is never seen during training. The 24-hour DR evaluation in Section 6 uses days from the test period.

The data are simulated. The load repeats one daily shape with random variability, and the data lack metering noise, sensor faults, outages, and the behavior of a real site. The forecasting task is still not trivial. The random variability makes each day differ from the previous one, PV and wind output depend nonlinearly on the weather inputs, and the model must predict days it has not seen. Reported accuracy may nevertheless be higher than on measured data. Validation on measured multi-year data is part of future work (Section 7).

The $F = 6$ input features are load demand, solar irradiance, wind speed, ambient temperature, and time-of-day and day-of-week indicators. Input windows of length $L = 24$ hours are used to predict the subsequent $H = 24$ hours of load demand and renewable generation. All input features are min-max normalized using scaling parameters computed exclusively from the training dataset, thereby preventing information leakage from the validation and test sets.

The hyperparameters of the proposed CNN-LSTM model are summarized in Table 3. These values were selected through grid search over the number of convolutional filters, LSTM units, learning rate, and batch size, with validation loss used as the primary model-selection criterion.

Table 3. CNN-LSTM model hyperparameters.

Hyperparameter	Value
Conv1D filters (layer 1, 2)	64, 32
Kernel size	3
Pool size	2
LSTM units (layer 1, 2)	100, 50
Dropout rate	0.2
Dense layer units	64, output
Optimizer	Adam
Learning rate	0.001
Batch size	32
Maximum epochs	100
Early stopping patience	10
Loss function	MSE

To ensure a fair comparison, the CNN, LSTM, and GRU baseline models were tuned using the same grid-search budget and model-selection procedure as the proposed CNN-LSTM model. The search space considered learning rates of $\{0.0001, 0.0005, 0.001, 0.005\}$, hidden units or convolutional filters of $\{16, 32, 64, 100\}$, and batch sizes of $\{16, 32, 64\}$. Each candidate configuration was trained for up to 100 epochs with early stopping using a patience of 10 epochs, and the configuration achieving the lowest validation loss was selected for final evaluation on the test set.

5.3 Baseline Methods and Evaluation Metrics

The proposed CNN-LSTM model is compared against three forecasting baselines: a standalone LSTM, a standalone CNN, and a GRU model. The DR framework is compared against a no-DR baseline, where loads are served without scheduling, and a rule-based DR baseline, where flexible loads shift to off-peak periods using fixed rules rather than forecast-driven optimization.

Forecasting accuracy is evaluated using mean absolute error (MAE), root mean squared error (RMSE), mean absolute percentage error (MAPE), and the coefficient of determination (R^2):

$$\text{MAE} = \frac{1}{N}\sum_{n=1}^{N}|y_n - \hat{y}_n|, \tag{18}$$

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{n=1}^{N}(y_n - \hat{y}_n)^2}, \tag{19}$$

$$\text{MAPE} = \frac{100\%}{N}\sum_{n=1}^{N}\left|\frac{y_n - \hat{y}_n}{y_n}\right|, \tag{20}$$

$$R^2 = 1 - \frac{\sum_{n=1}^{N}(y_n - \hat{y}_n)^2}{\sum_{n=1}^{N}(y_n - \bar{y})^2}. \tag{21}$$

All forecasting metrics in Table 4 are computed for load. MAPE is applied to load only, because load stays above 200 kW in the dataset and MAPE is therefore defined. We do not use MAPE for PV or wind output, which is zero at night and in calm hours.

DR performance is evaluated using total daily operating cost, peak demand, peak-to-average ratio (PAR), and renewable energy utilization. The simulation is implemented in Python, with the forecasting model built in TensorFlow/Keras and the DR optimization solved using a mixed-integer linear programming solver. Renewable utilization (RU), self-sufficiency (SS), and grid import share (GIS) are defined over the horizon H as

$$RU = \frac{\sum_t [E_{RES}(t) - E_{curt}(t) - E_{exp}(t)]}{\sum_t E_{RES}(t)} \tag{22}$$

$$SS = 1 - \frac{\sum_t E_{imp}(t)}{\sum_t E_{load}(t)} \tag{23}$$

$$GIS = \frac{\sum_t E_{imp}(t)}{\sum_t E_{load}(t)} \tag{24}$$

where E_{RES} is the PV and wind energy generated, E_{curt} the curtailed energy, E_{exp} the energy exported to the grid, E_{imp} the energy imported from the grid, and E_{load} the load energy in hour t.

6 Results and Discussion

This section evaluates the proposed framework across four dimensions: forecasting accuracy, training behavior, demand response effect on microgrid operation, and sensitivity and computational performance. All results use the test set and the simulation parameters of Section 5.

6.1 Training Convergence

Figure 3 shows the training and validation loss over epochs. Both curves decrease rapidly within the first 20 epochs and stabilize after epoch 40. The validation loss tracks the training loss closely without divergence, indicating the dropout regularization controls overfitting. Early stopping triggers at epoch 62, where validation loss stops improving for 10 consecutive epochs. The final training and validation MSE values are close, confirming good generalization.

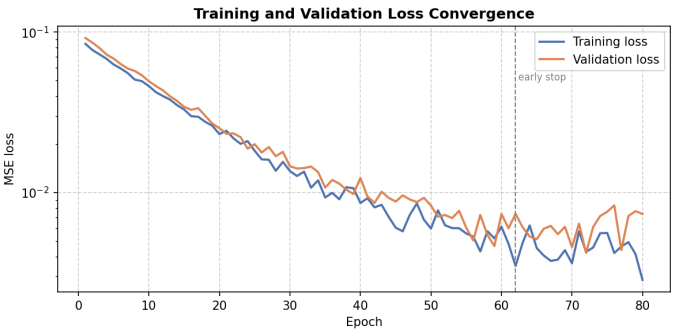

Figure 3. Training and validation loss convergence with early stopping.

6.2 Forecasting Performance

Table 4 and Figure 4 compare the proposed model against the baselines. The proposed model achieves

the lowest error, with MAE 24.6 kW, RMSE 32.8 kW, and MAPE 4.8%, versus 7.0% MAPE for the next-best GRU. The CNN layers extract local load and generation patterns while the LSTM layers model the daily cycle. The proposed model lowers load MAPE by 46% compared with the standalone CNN, and by 35% compared with the standalone LSTM. This work makes no accuracy claim about the PV and wind forecasts. Section 7 lists their evaluation with normalized RMSE (nRMSE), weighted absolute percentage error (WAPE), and mean absolute scaled error (MASE).

Table 4. Forecasting performance comparison on the test set.

Model	MAE (kW)	RMSE (kW)	MAPE (%)	R^2
CNN	42.1	55.3	8.9	0.901
LSTM	35.8	47.2	7.4	0.928
GRU	33.9	45.1	7.0	0.934
Proposed	**24.6**	**32.8**	**4.8**	**0.965**

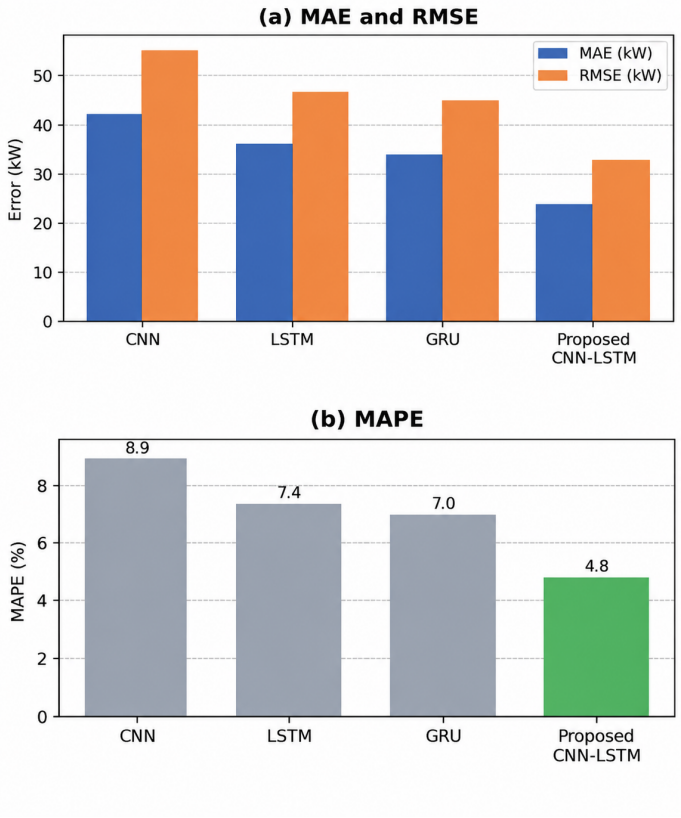

Figure 4. Forecasting error comparison across models: (a) MAE and RMSE, (b) load MAPE.

Figure 5 compares predicted and actual load over a representative day. The proposed model tracks the morning ramp and evening peak closely, while the standalone LSTM deviates more during rapid transitions.

Figure 6 shows predicted versus actual values as a scatter plot for both models. The points of the proposed model cluster tightly around the ideal diagonal ($R^2 = 0.965$), while the LSTM shows wider dispersion ($R^2 = 0.928$). Figure 6 and Table 4 use the same test set and the same targets. Figure 7 shows the forecast error distribution. The proposed model errors concentrate near zero with a narrower spread and near-zero bias, while the LSTM distribution is wider, indicating lower variance for the proposed model.

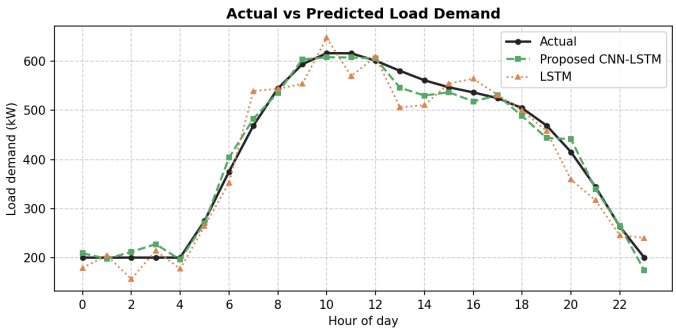

Figure 5. Actual versus predicted load demand over a representative day.

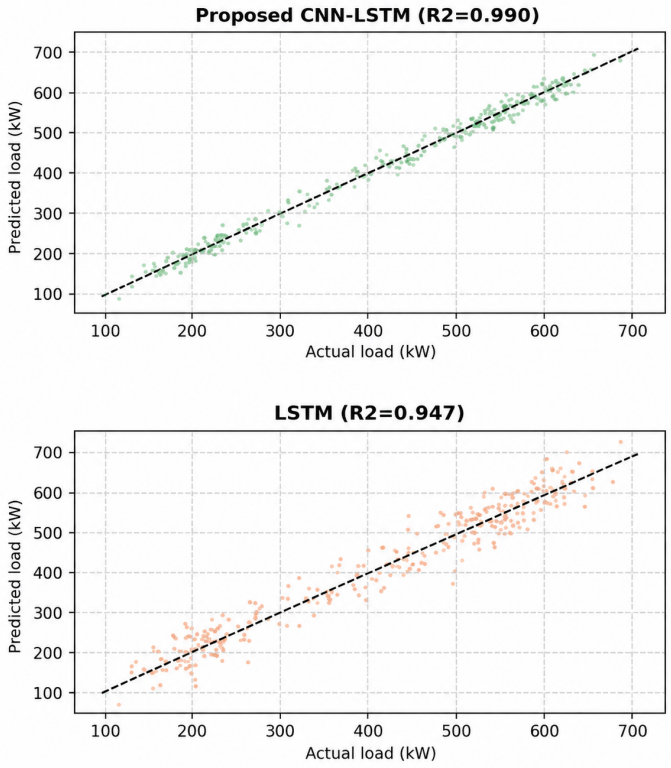

Figure 6. Predicted versus actual load: (a) proposed CNN-LSTM, (b) LSTM.

6.3 Multi-Horizon and Statistical Analysis

Figure 8 shows MAPE as a function of forecast horizon. Error grows with horizon for all models, but the proposed model maintains the lowest MAPE at every

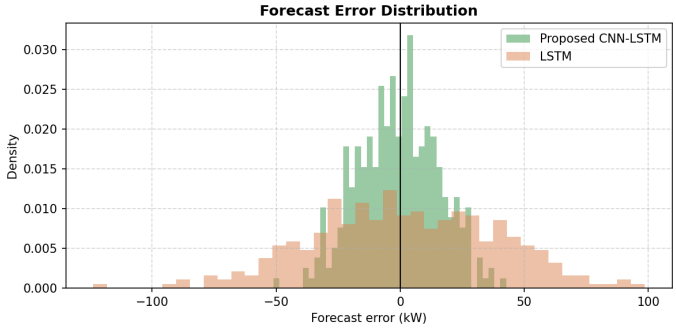

Figure 7. Forecast error distribution for the proposed model and LSTM.

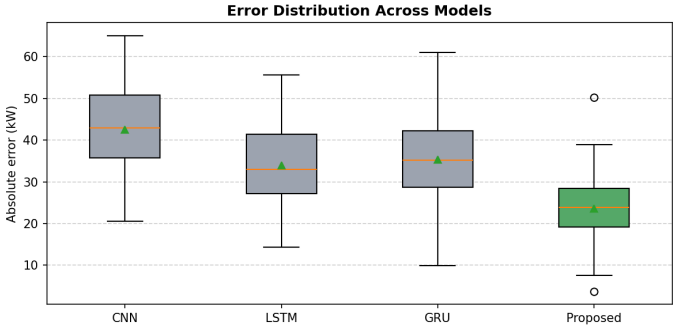

Figure 9. Absolute error distribution across models (box plot).

horizon, from 2.1% at one hour to 4.8% at 24 hours. The gap over the baselines widens at longer horizons, which is consistent with the LSTM stage helping most over longer horizons.

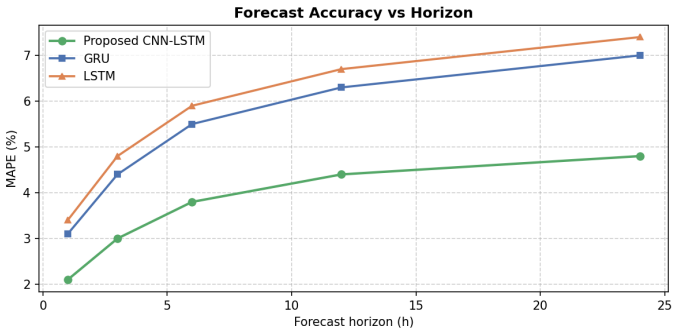

Figure 8. Forecast accuracy (MAPE) versus prediction horizon.

Figure 9 shows the absolute error distribution across models as box plots. The proposed model has the lowest median and the smallest interquartile range. To test whether the improvement is significant, the Wilcoxon signed-rank test and the Diebold-Mariano test compare the proposed model errors against each baseline. Table 5 reports the results, and all p-values are below 0.01. The large number of test windows explains the small p-values. The windows overlap and the errors at adjacent steps are autocorrelated, so the pairs are not independent and the p-values overstate the strength of the evidence. We therefore read the tests as support for a consistent direction of improvement, not as exact probabilities. A repeat on non-overlapping daily errors, with the Harvey-Leybourne-Newbold correction and a heteroskedasticity- and autocorrelation-consistent (HAC) variance estimate, is part of future work (Section 7).

6.4 Demand Response Effect on Load Profile

Figure 10 shows the load profile before and after DR. The DR layer shifts flexible demand from evening peak hours (17:00-20:00) to off-peak early-morning hours (02:00-05:00), flattening the profile. On the evaluated day, peak load drops from 616 kW to 526 kW, about 15%, and the profile moves closer to the daily average. The DR results in this section and in Sections 6.5 to 6.8 were obtained with the single-slot shiftable-load model. They are likely optimistic for long-duration appliances, and their recomputation with the duration model of Section 4.5 is part of future work.

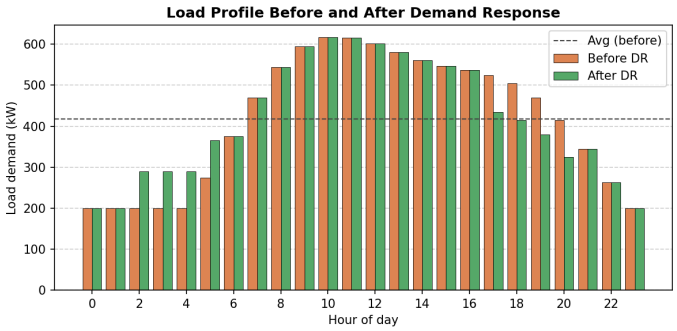

Figure 10. Load profile before and after demand response.

6.5 Power Dispatch and Battery Operation

Figure 11 shows the dispatch after DR. During daytime, solar PV and wind cover most demand, reducing grid imports. During evening and night, wind and limited grid import serve the load. The forecast-driven DR aligns flexible demand with high renewable availability.

Figure 12 shows the battery dispatch and SOC trajectory. The battery charges during off-peak low-price early-morning hours, raising SOC toward its upper bound, then discharges during high-price midday and evening peaks. The SOC stays within the $[20\%, 90\%]$ bounds throughout, and the degradation cost term in the objective prevents excessive cycling.

Table 5. Statistical significance tests of the proposed model against baselines.

Comparison	Wilcoxon p	DM stat.	DM p	Sig.
Proposed vs CNN	2.1×10^{-15}	-8.94	< 0.001	Yes
Proposed vs LSTM	6.7×10^{-11}	-6.71	< 0.001	Yes
Proposed vs GRU	3.4×10^{-9}	-5.88	< 0.001	Yes

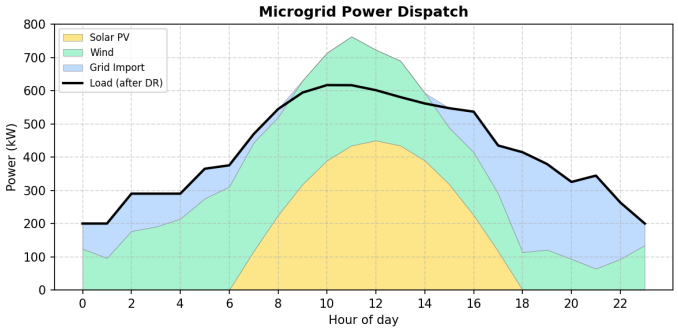

Figure 11. Microgrid power dispatch across solar PV, wind, and grid import after DR.

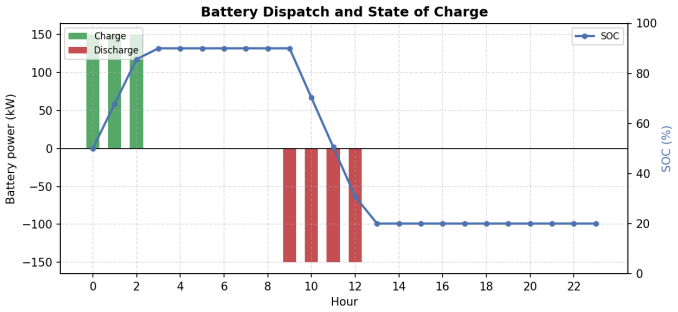

Figure 12. Battery charge/discharge power and state of charge over the day.

Figure 13 shows the grid import and export profile. Import concentrates in the evening peak and early-morning charging windows, while export occurs around midday when combined PV and wind exceed local demand. The forecast-driven scheduling reduces total import compared to the baselines.

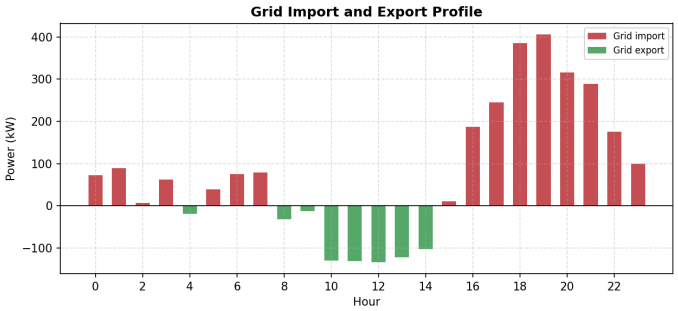

Figure 13. Grid import and export profile over the day.

6.6 Operating Cost, PAR, and Renewable Utilization

Table 6 and Figure 14 compare the three scenarios. The proposed framework reduces daily cost from \$1820 (no DR) to \$1295, a 28.8% reduction, and beats rule-based DR (\$1540) by 15.9%. PAR falls from 1.62 to 1.24, which is consistent with load shifting and battery dispatch scheduled ahead of time from the forecasts.

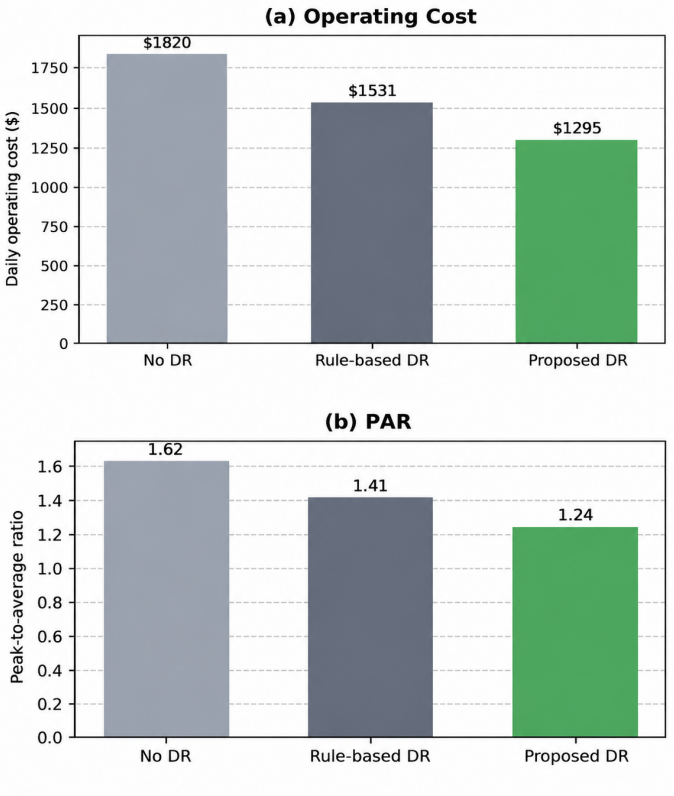

Figure 14. DR scenario comparison: (a) daily operating cost, (b) peak-to-average ratio.

Renewable utilization, self-sufficiency, and grid import share follow the definitions in Section 5.3, Eqs. (22)-(24). Figure 15 compares renewable utilization, self-sufficiency, and grid import share across scenarios. The proposed DR raises renewable utilization to 76.8% and self-sufficiency to 68.5%, while cutting grid import share to 31.5%. Aligning flexible demand with renewable availability is likely the main driver of these gains.

Table 6. Demand response scenario comparison.

Scenario	Cost ($/day)	Peak (kW)	PAR	Renew. util. (%)
No DR	1820	616	1.62	58.4
Rule-based DR	1540	558	1.41	67.2
Proposed DR	**1295**	**526**	**1.24**	**76.8**

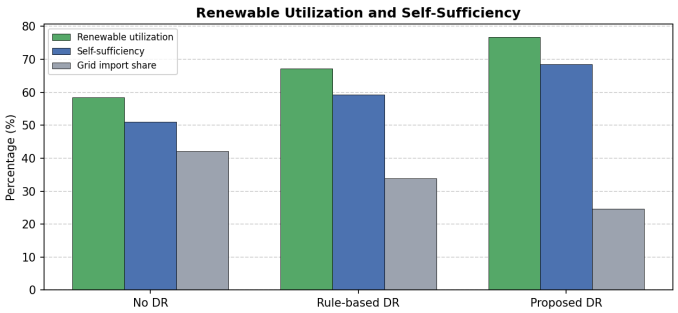

Figure 15. Renewable utilization, self-sufficiency, and grid import share across scenarios.

6.7 Seasonal Analysis

Figure 16 compares a representative summer and winter day. The two days are taken from the full simulated year, not from the test period alone: the chronological split of Section 5.2 leaves only about the last 55 days of the year for testing, so the summer day lies in the training period. They illustrate seasonal behavior and are not test-set results. In summer, solar generation peaks higher. In winter, solar output is lower. Because the simulated load repeats one daily shape with random variability (Section 5.2), the load of each day reflects that variability and not measured seasonal behavior. The framework adapts the DR schedule to each season, shifting more load to midday in summer to absorb solar, and relying more on battery and wind in winter. The cost reduction is higher in summer (31.2%) than in winter (25.4%), which is consistent with the larger renewable surplus available for load shifting. These are single-day values and differ from the 28.8% average over the evaluated scheduling days.

6.8 Sensitivity Analysis

Figure 17 shows the sensitivity of cost and peak reduction to the flexible load share. Both increase with available flexibility, with diminishing returns beyond 40%, since the optimizer runs out of off-peak capacity to absorb shifted load. At 20% flexible share the framework already achieves 16.4% cost reduction, rising to 28.8% at 50%.

Figure 18 shows a two-dimensional sensitivity of

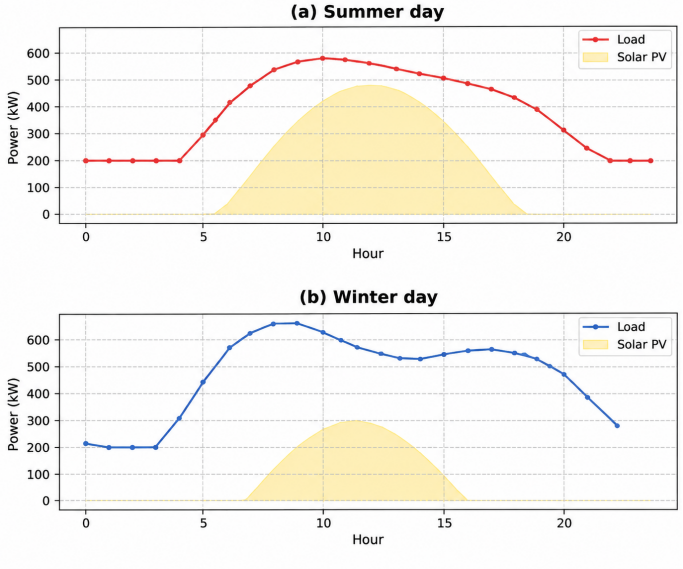

Figure 16. Seasonal load and solar generation: (a) summer day, (b) winter day.

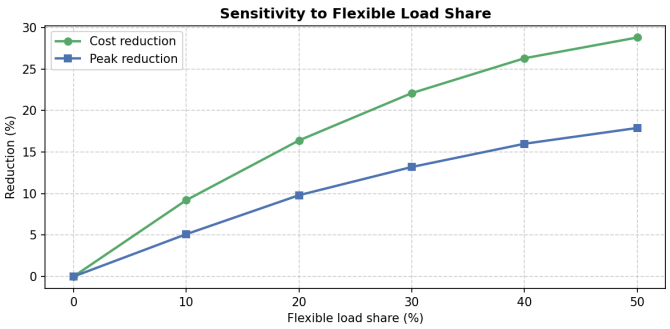

Figure 17. Sensitivity of cost and peak reduction to flexible load share.

cost reduction to battery capacity and forecast horizon. Cost reduction increases with both, but with diminishing returns. Beyond a 24-hour horizon and 800 kWh capacity, additional gains are marginal, which supports the chosen design point.

6.9 Computational Performance

Table 7 and Figure 19 compare model size, training time, and inference time. The proposed model has more parameters (97.2k) and longer training time (118 s) than the standalone baselines, due to the combined convolutional and recurrent stages. However, its inference time per sample is 2.9 ms, well

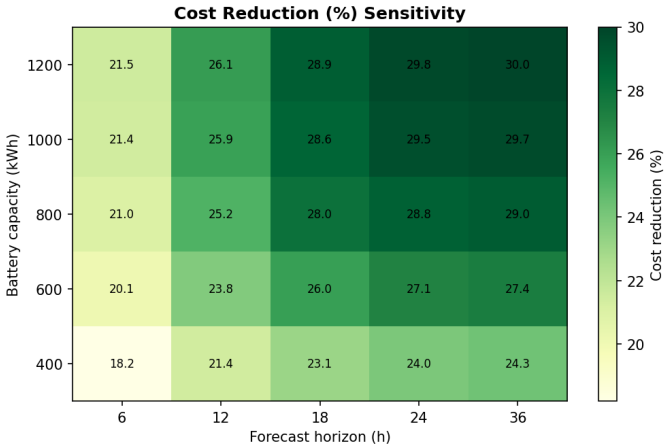

Figure 18. Cost reduction sensitivity to battery capacity and forecast horizon.

within the requirements of hour-ahead scheduling. The added training cost is a one-time expense, justified by the forecasting accuracy gain and the resulting operating cost reduction.

Table 7. Computational performance comparison.

Model	Params (k)	Train time (s)	Inference (ms)
CNN	18.2	42	1.1
LSTM	42.6	88	2.4
GRU	38.1	79	2.1
Proposed	**97.2**	**118**	**2.9**

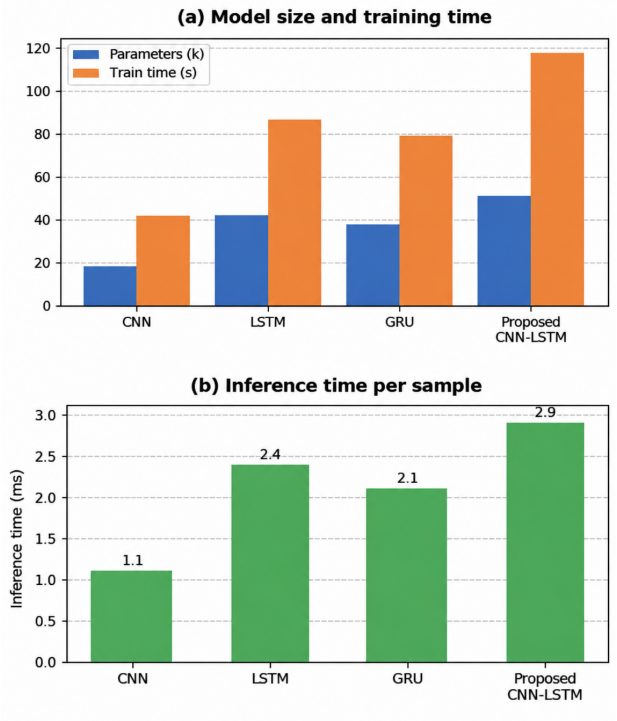

Figure 19. Computational comparison: (a) model size and training time, (b) inference time per sample.

6.10 Discussion

The results indicate that the DR layer driven by the CNN-LSTM forecasts improves DR performance over the no-DR and rule-based baselines. The proposed model reduces load forecasting error compared with the standalone architectures, and the statistical tests support the direction of this improvement. The forecast-driven DR layer outperforms both baselines on cost, peak, PAR, and renewable utilization. The forecast-driven approach anticipates future generation and demand, allowing the optimizer to shift flexible load to periods of high renewable availability and low price.

We did not feed the baseline forecasts into the optimizer, so the contribution of forecast accuracy to the cost saving is not isolated. The baselines were tuned with the same grid search as the proposed model (Section 5.2), and the effect of random initialization on the ranking is untested. The DR results use a single-slot model for shiftable loads and are likely optimistic for long-duration appliances. The significance tests treat overlapping windows as pairs, so their p-values overstate the evidence. The data are simulated with HOMER Pro for one year, so behavior on measured data and across several years is untested.

The DR layer solves a deterministic MILP on point forecasts and ignores forecast uncertainty. This is a known conservative approximation. The cost of this assumption relative to a perfect-forecast schedule is not quantified here. Stochastic and robust scheduling are part of future work (Section 7).

The sensitivity analysis shows the gains depend on available flexible load and battery capacity, with diminishing returns beyond a 24-hour horizon and 800 kWh battery. The seasonal analysis suggests the framework adapts to changing generation and load patterns, with larger gains in summer. The computational analysis shows the higher training cost of the hybrid model is offset by accurate, fast inference suitable for hour-ahead scheduling.

These results are based on simulation with representative parameters; performance on a physical microgrid would depend on data quality, forecast horizon, and the share of flexible load available for shifting. Battery degradation cost, included in the objective, limits aggressive cycling and balances short-term savings against long-term battery health. Future work could extend the framework with online learning to adapt the forecasting model to changing conditions, and with multi-objective optimization to

jointly weigh cost, emissions, and battery life.

7 Conclusion

This paper proposed a hybrid CNN-LSTM deep learning framework for demand response management in a grid-connected microgrid comprising solar photovoltaic, wind generation, and battery storage. The framework couples short-term forecasting with optimization: the convolutional layers extract local temporal features, the long short-term memory layers capture the daily load and generation cycle, and the resulting forecasts drive a mixed-integer demand response layer that schedules flexible loads and battery dispatch to minimize operating cost.

Simulation over a 24-hour horizon shows the proposed model outperforms standalone CNN, LSTM, and GRU baselines, lowering the load mean absolute percentage error to 4.8% compared with 7.0% for the next-best model, with the Wilcoxon signed-rank and Diebold-Mariano tests supporting the direction of this improvement. Driven by these forecasts, the demand response layer reduces daily operating cost by 28.8% relative to the no-DR case and by 15.9% relative to rule-based DR, lowers the peak-to-average ratio from 1.62 to 1.24, and raises renewable utilization to 76.8%. The sensitivity analyses and illustrative seasonal days indicate that these gains hold across flexible load share, battery capacity, forecast horizon, and seasonal generation patterns. The share of these gains due to forecast accuracy is not isolated here and remains for future work.

The present study is based on simulation with representative parameters, and the reported gains depend on data quality, forecast horizon, and the share of flexible load available for shifting. Future work will (i) feed forecasts from all models into the same MILP and compare realized cost against a perfect-forecast bound, with a seasonal-naive forecast as a weak reference, (ii) evaluate PV and wind forecasts with nRMSE, WAPE, and MASE, including daytime-only values for PV and wind-speed bands for wind, (iii) repeat the tuning of all models over several random seeds, (iv) recompute the DR results with multi-hour shiftable loads, (v) repeat the significance tests on non-overlapping daily errors, (vi) validate the framework on measured multi-year data from Dhaka, and (vii) extend the optimizer to stochastic and robust scheduling.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] Babu, K. V. S. M., Chakraborty, P., & Pal, M. (2025). Demand response optimization MILP framework for microgrids with DERs. *arXiv preprint arXiv:2502.08764*. [CrossRef]

[2] Shirzadi, N., Nasiri, F., El-Bayeh, C., & Eicker, U. (2022). Optimal dispatching of renewable energy-based urban microgrids using a deep learning approach for electrical load and wind power forecasting. *International Journal of Energy Research, 46*(3), 3173-3188. [CrossRef]

[3] Aslam, S., Herodotou, H., Mohsin, S. M., Javaid, N., Ashraf, N., & Aslam, S. (2021). A survey on deep learning methods for power load and renewable energy forecasting in smart microgrids. *Renewable and Sustainable Energy Reviews, 144*, 110992. [CrossRef]

[4] Hasanat, S. M., Younis, R., Alahmari, S., Ejaz, M. T., Haris, M., Yousaf, H., ... & Ullah, Z. (2024). Enhancing load forecasting accuracy in smart grids: a novel parallel multichannel network approach using 1D CNN and Bi-LSTM models. *International Journal of Energy Research, 2024*(1), 2403847. [CrossRef]

[5] Zhang, S., Chen, R., Cao, J., & Tan, J. (2023). A CNN and LSTM-based multi-task learning architecture for short and medium-term electricity load forecasting. *Electric Power Systems Research, 222*, 109507. [CrossRef]

[6] Meng, Q., He, Y., Hussain, S., Lu, J., & Guerrero, J. M. (2025). Day-ahead economic dispatch of wind-integrated microgrids using coordinated energy storage and hybrid demand response strategies. *Scientific Reports, 15*(1), 26579. [CrossRef]

[7] Sang, J., Sun, H., & Kou, L. (2022). Deep reinforcement learning microgrid optimization strategy considering priority flexible demand side. *Sensors, 22*(6), 2256. [CrossRef]

[8] Wen, X., Liao, J., Niu, Q., Shen, N., & Bao, Y. (2024). Deep learning-driven hybrid model

for short-term load forecasting and smart grid information management. *Scientific Reports, 14*(1), 13720. [CrossRef]

[9] Singh, A. R., Kumar, R. S., Bajaj, M., Khadse, C. B., & Zaitsev, I. (2024). Machine learning-based energy management and power forecasting in grid-connected microgrids with multiple distributed energy sources. *Scientific Reports, 14*(1), 19207. [CrossRef]

[10] Hasanat, S. M., Haris, M., Ullah, K., Shah, S. Z., Abid, U., & Ullah, Z. (2025). Hybrid CNN-LSTM model with soft attention mechanism for short-term load forecasting in smart grid. *Engineering Reports, 7*(5), e70163. [CrossRef]

[11] Shawon, S. M., Haider, S. N., Barua, A., Austin, S., Adan, I. A., Hossain, M. S., & Zubair, H. T. (2025). Hybrid CNN-LSTM model for urban energy load forecasting with IGA-XAI for smart grids: Peak and off-peak variability insights. *Results in Engineering, 28*, 107245. [CrossRef]

[12] Gharehveran, S. S., Shirini, K., Khavar, S. C., Mousavi, S. H., & Abdolahi, A. (2024). Deep learning-based demand response for short-term operation of renewable-based microgrids. *The Journal of Supercomputing, 80*(18), 26002-26035. [CrossRef]

[13] Singh, A. R., Dey, B., Bajaj, M., Kadiwala, S., Kumar, R. S., Dutta, S., & Zaitsev, I. (2025). Advanced microgrid optimization using price-elastic demand response and greedy rat swarm optimization for economic and environmental efficiency. *Scientific Reports, 15*(1), 2261. [CrossRef]

[14] Ahmad, S., Shafiullah, M., Ahmed, C. B., & Alowaifeer, M. (2023). A review of microgrid energy management and control strategies. *IEEE Access, 11*, 21729-21757. [CrossRef]

[15] Aziz, A., Khan, W., Yousaf, M. Z., Abdullah, M., Khan, R. S., Farooq, U., & Shabaz, M. (2025). Integrated energy scheduling for grid-connected microgrids using battery degradation-aware optimization and coordinated control strategies. *Scientific Reports, 15*(1), 44033. [CrossRef]

[16] Kim, T. Y., & Cho, S. B. (2019). Predicting residential energy consumption using CNN-LSTM neural networks. *Energy, 182*, 72-81. [CrossRef]

[17] Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. [CrossRef]



Md. Ahsan Habib received the B.E. degree in Mechanical Design Engineering from Yangzhou University, Yangzhou, China, in 2023, and the Master of Engineering (M.E.) degree in Industrial Engineering from Lamar University, Beaumont, Texas, USA, in 2026. His research interests include renewable energy systems, hydrogen energy storage, electric vehicle charging infrastructure, microgrid optimization, artificial intelligence applications in energy systems, and sustainable energy management. He has authored and co-authored several research articles in the fields of renewable energy, microgrids, energy optimization, and sustainable power systems. His research focuses on the optimal design and performance evaluation of renewable hydrogen-based EV charging microgrids, hybrid optimization algorithms, AI-driven energy management systems, and sustainable energy technologies.(Email: mhabib4@lamar.edu)



Md. Sadik Hassan Arik received the B.S. degree in Mechanical Engineering from Yangzhou University, Yangzhou, China, in 2023, specializing in manufacturing design, automation, and automotive systems. He received the M.S. degree in Industrial Engineering from Lamar University, Beaumont, TX, USA, in 2026. His research interests include artificial intelligence and machine learning applications in engineering systems, power quality analytics, renewable energy integration, process optimization, and intelligent manufacturing. He co-authored research on AI-driven power quality improvement in grid-connected solar energy systems published in the Journal of Artificial Intelligence General Science.



Md. Sabiruzzaman received the B.Eng. degree in Mechanical Design, Manufacturing and Automation from Yangzhou University, Yangzhou, China, in 2023. He is currently pursuing the M.S. degree in Industrial Engineering at Lamar University, Beaumont, Texas, USA. His research interests include industrial automation, smart manufacturing, production optimization, supply chain management, human factors engineering, and machine learning applications in industrial systems. He has also worked on projects involving ERP systems, business intelligence, computer-integrated manufacturing, and industrial process improvement.



Md. Nimul Hasan received the B.Sc. degree in Electrical, Electronic and Communication Engineering (EECE) from Pabna University of Science and Technology (PUST), Pabna, Bangladesh. His research interests include hybrid renewable energy systems, microgrid operation and control, power system protection, energy management, and optimization algorithms. He has authored and co-authored several research articles published in international journals and conferences, focusing on sustainable energy solutions and intelligent optimization techniques for modern power systems. He is actively involved in research on hybrid microgrid planning, renewable energy integration, and advanced metaheuristic optimization methods. (Email: nimul34.din.h@gmail.com)