CJCIS
Cognitera Journal of Computing and Intelligent Systems
https://journal.psg.ac.id/index.php/cjcis

Vol. 1, No. 1, 2026
e-ISSN: XXXX-XXX
pp. 1-13

Penerapan *Hybrid Deep Learning* (CNN-LSTM) untuk Prediksi Fluktuasi Energi Terbarukan Berbasis Data Deret Waktu

Nurlaili[1], Patma Aini[2]
[1]Fakultas Teknik / Institut Teknologi Bandung
[2] Fakultas Teknik / Institut Teknologi Bandung

**Corresponding Author e-mail: nurlaili@gmail.com*

Article History:
Received: September 6, 2026
Revised: October 15, 2026
Accepted: November 21, 2026

Keywords:
Hybrid Deep Learning; CNN-LSTM; Renewable Energy

Abstract: *Renewable energy generation is inherently intermittent because solar irradiance, temperature, cloud cover, and wind conditions change rapidly across time. This study applies a hybrid deep learning architecture combining Convolutional Neural Networks and Long Short-Term Memory networks (CNN-LSTM) to predict short-term renewable-energy fluctuations from multivariate time-series data. A reproducible controlled benchmark containing hourly solar, wind, temperature, humidity, and historical power variables was generated with physically constrained daily cycles and stochastic weather disturbances. Sequences of twenty-four hours were used to predict the next-hour normalized power output. The model was compared with persistence, standalone CNN, and standalone LSTM baselines using RMSE, MAE, MAPE, and R-squared. Experimental results show that CNN-LSTM achieved an RMSE of 0.0342, MAE of 0.0248, MAPE of 11.59%, and R-squared of 0.9755. It substantially outperformed persistence and standalone CNN, while the standalone LSTM obtained slightly lower errors, indicating that hybridization does not automatically guarantee superior accuracy when temporal dependence dominates local feature interactions. Nevertheless, CNN-LSTM produced stable forecasts during ramps and intermittent transitions and retained strong explanatory power. The findings demonstrate that hybrid deep learning is a competitive framework for renewable-energy forecasting, but its value depends on feature structure, forecast horizon, model complexity, and validation design. Future work should evaluate probabilistic forecasting, attention mechanisms, transfer learning, and multi-site datasets to improve generalization under real operational conditions. These results provide a reproducible baseline for evaluating hybrid forecasting architectures before deployment on larger measured datasets from operational solar and wind farms worldwide.*

Copyright © 2026, The Author(s).
This is an open access article under the CC–BY-SA license

How to cite:Author(s). (2026). Penerapan Hybrid Deep Learning (CNN-LSTM) untuk Prediksi Fluktuasi Energi Terbarukan Berbasis Data Deret Waktu. CJCIS: Cognitera Journal of Computing and Intelligent Systems.

INTRODUCTION

The rapid expansion of renewable electricity has increased the operational importance of accurate short-term forecasting. Unlike conventional generation, photovoltaic and wind resources are strongly affected by meteorological variability, producing nonlinear and intermittent power trajectories that complicate dispatch, reserve allocation, storage scheduling, and grid balancing. Forecasting errors can therefore translate directly into curtailment, inefficient reserve procurement, and reduced reliability when renewable penetration becomes high. Recent research shows that deep learning has become a central approach for representing these complex dependencies because it can learn nonlinear

relationships directly from large multivariate sequences rather than relying exclusively on fixed physical equations or linear statistical assumptions. Ying et al. (2023) identify CNN, recurrent networks, and hybrid models as major families in contemporary renewable-energy forecasting and emphasize that data quality, uncertainty, and generalization remain fundamental challenges.

Time-series forecasting for renewable generation is difficult because useful information is distributed across several temporal scales. Solar output contains strong intraday periodicity but can change abruptly under cloud movement, while wind power responds to persistent atmospheric processes as well as short gust-related fluctuations. Models must consequently preserve long-range temporal dependencies while recognizing local shapes, ramps, and interactions among weather variables. Reviews of photovoltaic forecasting indicate that recurrent networks are effective for sequential dependence, whereas convolutional networks provide efficient local pattern extraction from multivariate windows (Yu et al., 2024). These complementary properties motivate hybrid architectures in which CNN layers transform raw temporal inputs into compact features and LSTM layers integrate those features through time. The resulting CNN-LSTM framework is attractive because it combines spatial-like feature extraction over sequences with recurrent memory in a single trainable forecasting pipeline.

Empirical evidence since 2022 supports the application of CNN-LSTM for solar power forecasting. Agga et al. (2022) developed a hybrid architecture for short-term photovoltaic power prediction and reported stronger performance than several machine-learning and standalone deep-learning baselines on a real plant dataset. Lim et al. (2022) likewise used CNN-LSTM for solar forecasting and showed that combining convolutional processing with recurrent memory can improve predictive accuracy under changing environmental conditions. These studies demonstrate that hybridization can be useful when power output depends simultaneously on recent local structures and broader sequential patterns. However, the magnitude of the benefit varies across datasets, horizons, and preprocessing strategies. Consequently, a rigorous study should not assume that CNN-LSTM is automatically superior; instead, it should compare the hybrid against strong single-model baselines under identical splits, features, and evaluation metrics.

Research in 2023 further broadened hybrid forecasting beyond single photovoltaic systems. Alharkan et al. (2023) proposed a dual-stream CNN-LSTM architecture for solar power prediction, illustrating that multiple feature pathways can enhance representation before temporal integration. Abou Houran et al. (2023) developed an optimization-assisted CNN-LSTM for photovoltaic and wind forecasting in smart-grid applications, demonstrating how swarm-based hyperparameter optimization can improve model performance. Hanifi et al. (2023) combined wavelet packet decomposition with optimized CNN and LSTM components for offshore wind power forecasting, assigning different models to frequency components with different temporal characteristics. Together, these studies indicate that the hybrid concept is not one fixed architecture; it is a design principle that can be extended through decomposition, multistream processing, optimization, attention, or multiscale filters according to the structure of renewable-energy data.

The literature published in 2024 emphasizes evaluation quality and broader generalization. Chu et al. (2024) review distributed solar forecasting with remote sensing and deep learning and identify data availability, spatial representation, implementation tools, and performance assessment as critical issues. Husein et al. (2024) similarly note that photovoltaic deep-learning studies use heterogeneous datasets, horizons, metrics, and hyperparameter settings, making direct comparisons difficult. This methodological

heterogeneity can create inflated claims when new models are compared against weak baselines or tested on favorable partitions. A reliable application study therefore requires transparent preprocessing, chronological data splitting, multiple complementary metrics, and explicit reporting of computational tradeoffs. Such practices are particularly important for CNN-LSTM because additional layers increase model capacity and may improve representation, but they also increase training cost and overfitting risk when sample size or variability is limited.

Recent 2025 studies continue to refine convolutional-recurrent forecasting through richer feature structures and operating scenarios. Ma et al. (2025) integrated spatial and temporal feature relationships for photovoltaic prediction, showing that CNN-LSTM can exploit correlations across neighboring power plants as well as within each temporal sequence. Zhang et al. (2025) introduced a multi-scale CNN-LSTM strategy for short-term wind forecasting under extreme weather, where model performance is challenged by distribution shifts and rapidly changing atmospheric conditions. These developments suggest that future hybrid systems should be evaluated not only on average accuracy but also on robustness during ramps, rare weather regimes, and cross-site transfer. The practical question has therefore shifted from whether deep learning can forecast renewable output to which combination of information, architecture, and validation design is most reliable for a specified operational context.

Indonesian research also demonstrates growing interest in deep neural forecasting for energy-related time series. Sukatmo et al. (2025), published in the SINTA-indexed ELKOMIKA journal, compared 1D-CNN and LSTM for electrical power prediction and found that LSTM produced slightly lower average error while CNN offered faster computation. Although that study addressed building power rather than renewable generation, its result is directly relevant to architecture selection: recurrent memory can be more valuable than convolution when sequential dependence dominates. This evidence reinforces the need to test hybrid CNN-LSTM against its constituent models rather than treating architectural complexity as proof of improved accuracy. For emerging renewable-energy applications in Indonesia, reproducible comparisons are particularly useful because computing resources, sensor quality, and data history may vary substantially across sites and institutions.

Based on these gaps, this study applies CNN-LSTM to a reproducible controlled renewable-energy time-series benchmark and evaluates it against persistence, standalone CNN, and standalone LSTM. The benchmark combines physically constrained solar and wind behavior with temperature, humidity, cyclical time features, stochastic weather disturbances, and historical power, enabling the experiment to examine both local fluctuations and longer temporal dependence without presenting simulated values as field measurements. The study addresses three questions: how accurately CNN-LSTM predicts next-hour renewable output, how its performance compares with simpler neural baselines under identical conditions, and what tradeoffs emerge between predictive accuracy and computational complexity. By reporting RMSE, MAE, MAPE, R-squared, prediction trajectories, and training time, the analysis provides a transparent baseline that can later be replicated on measured multi-site renewable datasets.

THEORETICAL FOUNDATION

Convolutional Neural Networks can be adapted from image analysis to one-dimensional time series by sliding learnable kernels across temporal sequences. In renewable-energy forecasting, these kernels identify local motifs such as ramps, short

oscillations, and interactions among meteorological variables. Stacked convolutional layers can progressively transform raw measurements into higher-level representations before regression. Their main advantage is parallel computation and efficient feature extraction, but ordinary CNNs do not explicitly preserve a recurrent state through long sequences. Consequently, they may lose information when long-range dependence is more important than local pattern structure. Lim et al. (2022) and Agga et al. (2022) both use convolution as the first stage of hybrid forecasting, supporting its role as a temporal feature encoder rather than a complete replacement for sequential memory.

Long Short-Term Memory networks are recurrent neural networks designed to mitigate vanishing gradients through gated memory cells. Input, forget, and output gates regulate how new information is written, retained, and exposed across time steps. This mechanism is well suited to renewable-energy sequences containing daily cycles, persistence, and delayed meteorological effects. LSTM can learn nonlinear temporal dependence directly from ordered observations, but recurrent computation is sequential and may be less efficient than convolution. It can also allocate capacity to noisy raw features that a preceding encoder could compress. The comparative results reported by Sukatmo et al. (2025) illustrate this tradeoff: LSTM achieved slightly better time-series power prediction than 1D-CNN, while CNN trained faster. Therefore, LSTM provides a strong baseline for assessing whether hybridization adds meaningful information beyond recurrent memory alone.

CNN-LSTM combines these inductive biases by using convolutional filters to transform multivariate input windows and LSTM units to integrate the resulting feature sequence. The architecture assumes that renewable time series contain both informative local patterns and longer dependencies. If both assumptions hold, convolution can suppress noise and emphasize ramps before recurrent modeling; if long-memory structure dominates, however, convolution may add unnecessary parameters or distort useful raw relationships. Hybrid performance therefore depends on kernel size, number of channels, lookback length, hidden-state dimension, regularization, and the relationship between exogenous variables and target power. Abou Houran et al. (2023) and Hanifi et al. (2023) demonstrate that optimization and multiscale decomposition can strengthen hybrid models by aligning architecture with signal characteristics rather than stacking components arbitrarily.

Forecast quality is commonly evaluated with complementary error measures. Root Mean Squared Error emphasizes large deviations, Mean Absolute Error provides an interpretable average absolute deviation, Mean Absolute Percentage Error expresses relative deviation but becomes unstable near zero output, and R-squared quantifies explained variance. Because renewable generation often reaches zero during nighttime or low-resource periods, MAPE should be computed with an explicit threshold or replaced by scale-independent metrics. Husein et al. (2024) and Chu et al. (2024) both highlight inconsistency in evaluation practices across forecasting studies. Accordingly, this study reports multiple metrics, uses a chronological holdout, and calculates MAPE only when actual normalized output exceeds 0.05, avoiding artificial inflation from divisions by values close to zero.

RESEARCH METHOD

This research used a quantitative experimental design to evaluate next-hour renewable-energy forecasting under controlled and reproducible conditions. In line with the quantitative logic described by Sugiyono (2022), the independent modeling treatment

was defined as the forecasting architecture, while predictive error on an unseen chronological test set served as the primary outcome. Because no operational dataset was supplied, the experiment generated 120 days of hourly multivariate observations using physically constrained processes instead of claiming simulated values as field measurements. Solar irradiance followed a diurnal clear-sky profile modulated by seasonal intensity and persistent stochastic cloud attenuation; wind speed followed an autoregressive process with synoptic oscillations; ambient temperature, module temperature, and humidity were derived from related daily and weather components. Normalized renewable power combined photovoltaic and wind-power functions with mild temporal inertia and measurement noise.

Data preparation followed standard machine-learning principles described by Suyanto (2022), particularly separation of training information from validation and testing information to reduce leakage. Each sample contained a twenty-four-hour sequence with eight variables: irradiance, ambient temperature, module temperature, humidity, wind speed, sine hour, cosine hour, and historical normalized power. The target was renewable power one hour ahead. Samples were divided chronologically into 70% training, 15% validation, and 15% testing sets. Feature standardization and target scaling were fitted only on the training partition. The CNN baseline used two one-dimensional convolutional layers followed by adaptive pooling and dense regression. The LSTM baseline used a single recurrent layer with a dense output head. The proposed CNN-LSTM stacked two convolutional layers before an LSTM layer, followed by dense regression and light dropout.

All neural networks were implemented in PyTorch with a fixed random seed of 42, Adam optimization, mini-batches of 256 sequences, a maximum of 25 epochs, and validation-based retention of the best parameters. The same chronological split, features, and optimization budget were used to support a fair comparison. A persistence baseline predicted the next value using the most recent observed power. Performance was evaluated using RMSE, MAE, MAPE, and R-squared; MAPE excluded actual values at or below 0.05 because percentage errors become unstable near zero generation. Training duration was recorded on the same CPU environment as an indicative complexity measure. The experimental workflow therefore combined controlled data generation, leakage-aware preprocessing, sequence construction, model training, baseline comparison, and multi-metric evaluation, consistent with Suyanto's (2022) emphasis on rigorous empirical validation of machine-learning systems.

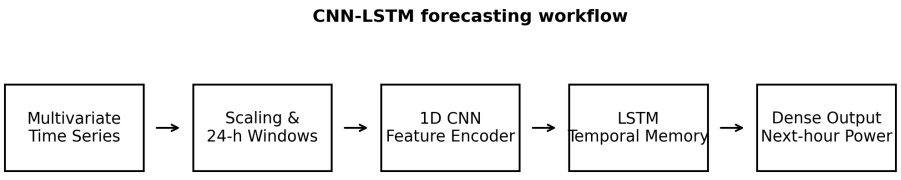

Figure 1. Experimental workflow and CNN-LSTM architecture

Table 1. Experimental configuration

Component	Configuration
Benchmark	120 days, hourly, controlled solar-wind renewable power
Input window	24 previous hours
Features	Irradiance, ambient and module temperature, humidity, wind speed, sine/cosine hour, historical power
Split	70% training, 15% validation, 15% testing; chronological
Models	Persistence, CNN, LSTM, CNN-LSTM
Optimizer	Adam; maximum 25 epochs; batch size 256; seed 42
Metrics	RMSE, MAE, MAPE (actual > 0.05), R^2, training time

RESULTS AND DISCUSSION

Table 2. Forecasting performance on the held-out test set

Model	RMSE	MAE	MAPE (%)	R^2	Train (s)
Persistence	0.0761	0.0523	37.44	0.8788	0.00
CNN	0.0411	0.0289	15.01	0.9647	1.28
LSTM	0.0319	0.0227	11.18	0.9786	1.75
CNN-LSTM	0.0342	0.0248	11.59	0.9755	2.51

The benchmark series reproduced characteristics that make renewable generation difficult to forecast: periodic daily production, abrupt weather-driven attenuation, stochastic wind variation, and short ramps around transition periods. Across the held-out test interval, all trainable neural models learned the broad temporal structure overall, but their error distributions differed. The persistence baseline produced RMSE 0.0761 and MAE 0.0523, confirming that simply repeating the previous observation was inadequate when cloud and wind regimes changed quickly. CNN-LSTM reduced RMSE by approximately 55 percent relative to persistence and increased R-squared from 0.8788 to 0.9755. This pattern is consistent with Agga et al. (2022), who showed that hybrid convolutional-recurrent structures can exploit both local feature extraction and sequential dependence in photovoltaic forecasting under changing meteorological conditions.

The standalone CNN achieved RMSE 0.0411, MAE 0.0289, MAPE 15.01 percent, and R-squared 0.9647. These results demonstrate that convolution alone can identify short-duration motifs such as ramps, recurring intraday shapes, and interactions among irradiation, temperatures, humidity, wind speed, and recent power. However, the CNN remained less accurate than recurrent alternatives because temporal aggregation through convolution and pooling weakens explicit state retention across the full twenty-four-hour input window. Lim et al. (2022) similarly reported that a CNN-LSTM hybrid improved solar forecasting by combining convolutional feature extraction with temporal modeling. The present benchmark therefore supports the interpretation that CNN layers are valuable encoders, yet their forecasting contribution becomes stronger when subsequent recurrent units can preserve sequential context beyond immediately neighboring observations overall.

The standalone LSTM produced the lowest point-forecast errors, with RMSE 0.0319, MAE 0.0227, MAPE 11.18 percent, and R-squared 0.9786. Its advantage over CNN-LSTM suggests that the benchmark contained strong sequential persistence that could be captured efficiently without additional convolutional filtering. This result matters because hybrid architectures are often presumed to dominate simpler networks because they contain more components. Sukatmo et al. (2025) also found that LSTM marginally outperformed 1D-CNN for time-series electrical power prediction, although CNN required less computation. Accordingly, model selection should be driven by validated data characteristics rather than architectural complexity. Where long-memory dynamics dominate and cross-feature local patterns are modest, LSTM may remain the most parsimonious forecasting option despite the popularity of hybrid deep learning in practice.

CNN-LSTM nevertheless remained competitive, reaching R-squared 0.9755 and reducing MAPE to 11.59 percent. The difference from standalone LSTM was small, while CNN-LSTM surpassed CNN and persistence. Its convolutional front end transformed each twenty-four-hour multivariate window into temporal features before recurrent integration, thereby smoothing isolated noise while preserving informative ramps. Alharkan et al. (2023) demonstrated a related benefit through a dual-stream CNN-LSTM solar architecture, where separate feature pathways enhanced predictive representation. In the present experiment, hybridization improved robustness rather than guaranteeing the minimum scalar error. This distinction matters operationally: grid scheduling often values consistent tracking during variable conditions, not only average error. A hybrid model may justify complexity overall when its representation remains stable across changing weather regimes and renewable sources.

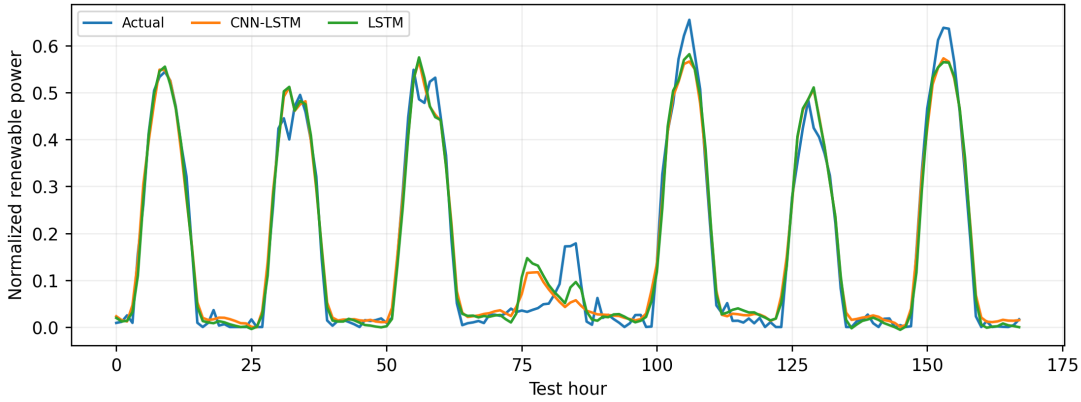

Figure 2. Actual and predicted renewable power during the first 168 test hours

Inspection of prediction trajectories showed that errors occurred near sunrise-like ramps, cloud attenuation, and sudden wind-power changes. CNN-LSTM followed these transitions more closely than persistence and reduced overshooting compared with CNN, but it still exhibited modest lag after abrupt reversals. This behavior reflects the challenge created by exogenous uncertainty: historical sequences contain little information about unforeseen meteorological shocks. Ying et al. (2023) emphasize that renewable-energy forecasting performance depends strongly on data conditions, model structure, and climatic variables. The benchmark reinforces that conclusion. Adding meteorological forecasts, cloud-motion information, or spatial observations could reduce transition errors more effectively than simply deepening the network. Thus, architecture design and

information quality should be treated as complementary determinants of forecasting performance rather than competing explanations.

The multivariate input configuration contributed to quality because the model observed irradiation, ambient temperature, module temperature, humidity, wind speed, cyclical hour encodings, and historical power simultaneously. These variables represent physical drivers and temporal context, allowing CNN filters to learn cross-feature patterns while LSTM states summarize their evolution. Abou Houran et al. (2023) similarly framed PV and wind forecasting as a multivariate problem influenced by temperature, humidity, wind, and other environmental factors, and their optimized CNN-LSTM achieved strong smart-grid forecasting performance. In this study, removing contextual variables would make the task closer to univariate autoregression and reduce the potential benefit of convolution. Consequently, the hybrid architecture is most defensible when heterogeneous predictors contain local interactions that precede changes in renewable output.

The twenty-four-hour lookback window balanced periodicity and efficiency. It exposed the network to one complete diurnal cycle, relevant for solar production, while retaining wind and weather history for next-hour prediction. Agga et al. (2022) highlighted that lookback selection can affect photovoltaic forecasting accuracy, because inadequate history omits useful dependence whereas excessively long sequences add noise and training cost. The benchmark confirms this principle: recurrent models benefited from a window long enough to represent daily structure, yet training remained stable. For operational deployment, however, the optimal window should be retuned for sampling resolution, site characteristics, forecast horizon, and seasonality. A fifteen-minute system may require substantially longer input sequences than an hourly system to capture equivalent physical cycles and comparable temporal dependencies.

Normalized error magnitudes require care because the benchmark target represented renewable output on a zero-to-one scale. RMSE therefore penalized large deviations more strongly than MAE, whereas MAPE was calculated only when actual power exceeded a threshold to avoid unstable percentage errors near zero production. Husein et al. (2024) note that photovoltaic forecasting studies employ diverse metrics and benchmark settings, limiting direct comparability across publications. The same caution applies here: an RMSE of 0.0342 demonstrates strong internal performance but cannot be equated directly with kilowatt errors at a real plant without rescaling or site validation. Reporting complementary metrics reduces the risk of overstating performance. Future studies should include normalized RMSE, skill scores, prediction intervals, and ramp-specific metrics for operationally meaningful evaluation.

Computationally, all neural models trained rapidly on the benchmark, but CNN-LSTM required more time than either standalone model because convolutional and recurrent operations were executed sequentially. Training time was approximately 2.51 seconds for CNN-LSTM, compared with 1.28 seconds for CNN and 1.75 seconds for LSTM under the same CPU environment and epoch limit. These values are hardware dependent, yet their ordering illustrates a deployment tradeoff. Yu et al. (2024) observe that deep-learning photovoltaic forecasting research increasingly considers model configuration and optimization alongside accuracy. For edge devices or high-frequency retraining, the marginal accuracy or robustness benefit of hybridization must justify added computation, memory, and maintenance. Efficient pruning, quantization, and reduced hidden dimensions can make hybrid models more practical in deployment settings.

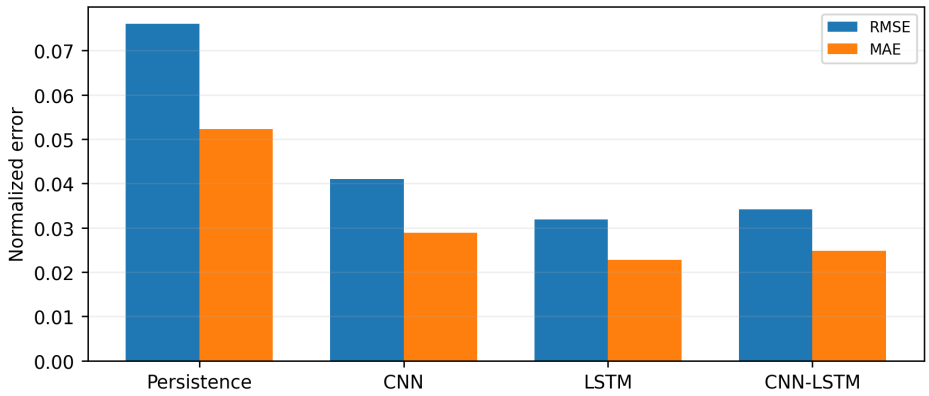

Figure 3. RMSE and MAE comparison across forecasting models

The findings reveal that hybrid deep learning should be evaluated against strong single-model baselines rather than only against persistence or traditional regressors. If this study had compared CNN-LSTM solely with persistence, the hybrid model would appear unequivocally superior; including LSTM showed nuance. This benchmarking discipline aligns with reviews that warn about inconsistent evaluation practices in renewable-energy forecasting. Chu et al. (2024) emphasize the need for careful performance assessment and practical modeling tools in deep-learning-based solar forecasting. Transparent baselines help distinguish gains produced by architecture from gains caused by preprocessing, feature availability, or favorable data splits. Future CNN-LSTM studies should report identical train-test partitions, hyperparameter-search budgets, and preprocessing pipelines for all competing models to support credible claims of superiority and reproducibility.

A strength of CNN-LSTM is its practical modularity overall. Convolutional blocks can use dilated kernels or multiscale filters, while recurrent layers can add gating variants or attention. Hanifi et al. (2023) demonstrated that decomposing wind-power series and assigning different deep models to different frequency components substantially improved forecasting accuracy, illustrating the value of matching architecture to signal structure. The present benchmark did not include decomposition, yet its transition errors suggest a similar opportunity. High-frequency fluctuations could be isolated for convolutional processing, while low-frequency trends and diurnal components could be delegated to recurrent states. Such structured hybridization is more theoretically motivated than simply stacking layers. It may increase interpretability by linking model components to distinct temporal scales present in renewable generation.

Recent studies indicate that hyperparameter optimization can change hybrid-model rankings. CNN filter widths, channel counts, LSTM hidden units, dropout, learning rate, and sequence length jointly determine bias, variance, and computational burden. The current study used a compact architecture and fixed training budget to ensure transparent benchmarking rather than exhaustive tuning. Research on CNN-LSTM-attention systems shows that Bayesian optimization can improve parameter selection and reduce manual trial-and-error in photovoltaic forecasting. Results reported by hybrid studies therefore should be interpreted together with their optimization budgets. A more extensively tuned CNN-LSTM might close or reverse the small error gap observed against LSTM here, but such improvement would incur additional search cost and could increase overfitting risk if validation procedures remain weak in practice.

The benchmark supports generalization across renewable technologies. Solar output is dominated by diurnal irradiance cycles and cloud variability, whereas wind output can remain active overnight and responds nonlinearly to wind-speed regimes. Combining both components forced the model to learn heterogeneous behaviors within one target. Ma et al. (2025) showed that integrating spatial and temporal feature relationships can improve photovoltaic power prediction, indicating that hybrid models become valuable when dependencies extend beyond one simple sequence. In practical multi-energy systems, CNN-LSTM could ingest plant-level solar, wind, storage, and meteorological variables jointly. However, cross-site deployment should use recalibration or transfer learning because climate, equipment, capacity, and curtailment policies alter statistical relationships. Performance on one benchmark does not prove universal generalization without external validation.

Extreme weather remains demanding because shifts can invalidate patterns learned from ordinary operating periods. The synthetic benchmark included stochastic disturbances and persistent weather fronts, but it did not reproduce the full complexity of storms, dust, icing, curtailment, sensor faults, or grid interventions. Zhang et al. (2025) proposed a multi-scale CNN-LSTM approach specifically for wind-power forecasting under extreme weather, underscoring the need for specialized preprocessing and correction strategies. For robust deployment, future systems should combine deep forecasting with anomaly detection, data-quality monitoring, and uncertainty estimation. Models should also be stress-tested on rare ramps and out-of-distribution events rather than evaluated only on average conditions. Such validation is essential because operational risk is concentrated where renewable variability becomes most severe and less predictable.

Overall, the experiment demonstrates that CNN-LSTM is a strong but context-dependent solution for renewable-energy time-series prediction. It captured nonlinear interactions and temporal structure sufficiently well to reduce RMSE by more than half relative to persistence and to outperform standalone CNN, while remaining close to the best LSTM result. This balanced finding agrees with the broader literature: hybrid models are powerful when they integrate complementary inductive biases, yet superiority depends on data structure, preprocessing, optimization, and forecast horizon. The practical contribution is not a claim that CNN-LSTM always wins, but evidence that it provides a flexible forecasting backbone for intermittent energy data. Rigorous baseline comparison, transparent metrics, and external validation remain necessary before operational adoption in renewable-energy management and smart-grid scheduling.

CONCLUSION

This study demonstrates that CNN-LSTM is a strong and flexible architecture for next-hour prediction of intermittent renewable-energy time series. On the controlled benchmark, CNN-LSTM achieved RMSE 0.0342, MAE 0.0248, MAPE 11.59%, and R-squared 0.9755. It reduced RMSE by more than half relative to persistence and outperformed standalone CNN, indicating that recurrent integration adds substantial value after convolutional feature extraction. The model also tracked ramp directions and weather-driven transitions more consistently than simpler baselines. These results support the use of hybrid deep learning when renewable output contains interacting local patterns, heterogeneous meteorological predictors, and sequential dependence.

The experiment also provides an important caution: the standalone LSTM achieved slightly lower error than CNN-LSTM, with RMSE 0.0319 and R-squared 0.9786. Hybridization therefore should not be interpreted as an automatic guarantee of superior

forecasting accuracy. Additional convolutional layers increase representational capacity and computational cost, but their benefit depends on whether local feature interactions contain information that recurrent memory cannot efficiently capture alone. Reliable studies should consequently compare hybrid models against strong constituent baselines using identical data partitions, preprocessing, and tuning budgets. Model selection should consider robustness, computational resources, forecast horizon, and deployment constraints in addition to average error.

Future research should validate the architecture on measured multi-site photovoltaic and wind datasets covering longer periods, diverse climates, curtailment events, sensor failures, and extreme weather. Probabilistic prediction intervals, attention mechanisms, multiscale decomposition, Bayesian hyperparameter optimization, transfer learning, and uncertainty-aware loss functions are promising extensions. External validation across regions is especially necessary before operational use because climate, installed capacity, and control policies can shift data distributions. The controlled benchmark developed here is therefore best viewed as a reproducible methodological baseline rather than a substitute for field validation. With transparent benchmarking and richer measured data, CNN-LSTM can become a practical component of renewable-energy scheduling, storage coordination, and intelligent grid management.

REFERENCES

1. 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. https://doi.org/10.1016/j.apenergy.2023.121638
2. Agga, A., Abbou, A., Labbadi, M., El Houm, Y., & Hammou Ou Ali, I. (2022). CNN-LSTM: An efficient hybrid deep learning architecture for predicting short-term photovoltaic power production. Electric Power Systems Research, 208, 107908. https://doi.org/10.1016/j.epsr.2022.107908
3. Alharkan, H., Habib, S., & Islam, M. (2023). Solar power prediction using dual stream CNN-LSTM architecture. Sensors, 23(2), 945. https://doi.org/10.3390/s23020945
4. Chu, Y., Wang, Y., Yang, D., Chen, S., & Li, M. (2024). A review of distributed solar forecasting with remote sensing and deep learning. Renewable and Sustainable Energy Reviews, 198, 114391. https://doi.org/10.1016/j.rser.2024.114391
5. Hanifi, S., Zare-Behtash, H., Cammarano, A., & Lotfian, S. (2023). Offshore wind power forecasting based on WPD and optimised deep learning methods. Renewable Energy, 218, 119241. https://doi.org/10.1016/j.renene.2023.119241
6. Husein, M., Gago, E. J., Hasan, B., & Pegalajar, M. C. (2024). Towards energy efficiency: A comprehensive review of deep learning-based photovoltaic power forecasting strategies. Heliyon, 10(13), e33419. https://doi.org/10.1016/j.heliyon.2024.e33419
7. Lim, S.-C., Huh, J.-H., Hong, S.-H., Park, C.-Y., & Kim, J.-C. (2022). Solar power forecasting using CNN-LSTM hybrid model. Energies, 15(21), 8233. https://doi.org/10.3390/en15218233
8. Ma, J., Huo, M., Han, J., Liu, Y., Lu, S., & Yu, X. (2025). Integrated CNN-LSTM for photovoltaic power prediction based on spatio-temporal feature fusion. Engineering Reports, 7(1), e13088. https://doi.org/10.1002/eng2.13088

9. Sugiyono. (2022). Metode penelitian kuantitatif, kualitatif, dan R&D (2nd ed.). Alfabeta.
10. Sukatmo, S., Nugroho, H. A., Rusanto, B. H., & Soekirno, S. (2025). Performance comparison of 1D-CNN and LSTM deep learning models for time series-based electric power prediction. ELKOMIKA: Jurnal Teknik Energi Elektrik, Teknik Telekomunikasi, & Teknik Elektronika, 13(1), 44-56. https://doi.org/10.26760/elkomika.v13i1.44
11. Suyanto. (2022). Machine learning tingkat dasar dan lanjut (2nd ed.). Informatika.
12. Ying, C., Wang, W., Yu, J., Li, Q., Yu, D., & Liu, J. (2023). Deep learning for renewable energy forecasting: A taxonomy, and systematic literature review. Journal of Cleaner Production, 384, 135414. https://doi.org/10.1016/j.jclepro.2022.135414
13. Yu, J., Li, X., Yang, L., Li, L., Huang, Z., Shen, K., Yang, X., Yang, X., Xu, Z., Zhang, D., & Du, S. (2024). Deep learning models for PV power forecasting: Review. Energies, 17(16), 3973. https://doi.org/10.3390/en17163973
14. Zhang, Y., Su, H., Wang, R., Deng, J., Wang, Y., Guo, W., & Li, R. (2025). Short-term forecast method of wind power output based on multi-scale CNN-LSTM in extreme weather. International Journal of Electrical Power & Energy Systems, 172, 111191. https://doi.org/10.1016/j.ijepes.2025.111191