www.ijarp.com ISSN 2456-9992 Page: 01-17

Research Article

International Journal Advanced Research Publications

Volume: 02
Issue: 11

FEDERATED AND DISTRIBUTED MACHINE LEARNING FOR RENEWABLE ENERGY FORECASTING AND PROSUMER ANALYTICS

***S. Swethanjali, S. Tarun**

Dept. of EEE, GMRIT Deemed to Be University, Rajam, Vizianagaram, Andhra Pradesh, India.

Article Received: 12 October 2026, Article Revised: 01 November 2026, Published on: 20 November 2026

***Corresponding Author: S. Swethanjali**

Dept. of EEE, GMRIT Deemed to Be University, Rajam, Vizianagaram, Andhra Pradesh, India.

Doi: https://doi-doi.org/101555/ijarp.9013

ABSTRACT

The growing deployment of renewable energy has increased the need for forecasting methods that are accurate, scalable, and respectful of data privacy. Federated and distributed machine learning offer a way for geographically separated energy resources to train shared models without transferring their raw datasets to a central location. This paper reviews the use of these approaches in renewable-energy forecasting and prosumer analytics. It considers solar and wind power prediction, smart-grid operation, energy management, demand forecasting, and privacy-aware analytics. The discussion also addresses practical issues, including non-IID data, communication overhead, cybersecurity, scalability, and differences in client availability and computing capacity. Based on the reviewed studies, a federated architecture is outlined in which PV plants, wind farms, buildings, battery systems, EV charging stations, and household prosumers act as distributed clients. Each client preprocesses its data and trains a local model, while a coordinator aggregate protected model updates. The reviewed work suggests that federated learning can support collaborative energy analytics while allowing participating organizations and users to retain control of their raw data.

KEYWORDS: Renewable Energy Forecasting, Federated Learning, Distributed Machine Learning, Prosumer Analytics, Privacy Preservation, Smart Grid, Deep Learning.

I. INTRODUCTION

A. Background and Motivation

The growing use of renewable sources, particularly solar photovoltaic (PV) and wind power, is changing the way electrical power systems are planned and operated. These sources can reduce dependence on fossil-fuel-based generation, but their output varies with weather and environmental conditions. Reliable forecasts are therefore important for power-system planning, unit scheduling, energy management, storage operation, and demand response [5]–[8]. Machine-learning (ML) and deep-learning (DL) methods are increasingly used for this purpose because they can learn nonlinear and time-dependent relationships from historical generation and meteorological data [5], [7].Many conventional forecasting systems collect measurements at a central location and train a model on the combined dataset. Although this arrangement gives the model access to data from several sites, it can be difficult to implement when measurements belong to different plants, organizations, buildings, or individual consumers. Energy records may reveal operating conditions, consumption habits, equipment information, or commercially sensitive details. Transferring all such records to one server can therefore raise privacy, security, ownership, and communication concerns [2]–[4].

Distributed machine learning and federated learning provide alternatives to collecting all data in one place. In federated learning, participating clients train a shared model using their own local datasets. Rather than sending raw measurements to the coordinator, they transmit model parameters or updates for aggregation [3]. Liu et al. describe this approach as useful when data and computing resources are spread across different locations and cannot readily be combined because of organizational, technical, or regulatory constraints [3]. This makes federated learning relevant to energy systems, where data are naturally distributed across locations and ownership boundaries.

B. Renewable-Energy Forecasting

Renewable-energy forecasting estimates future power generation from historical output, weather information, environmental measurements, and operating conditions. Solar and wind generation respond to different physical factors. Solar output is affected by irradiance, cloud cover, temperature, and atmospheric conditions, whereas wind generation depends on wind speed and direction, atmospheric stability, turbulence, and local geography [6].

ML and DL methods are widely studied for renewable-energy forecasting because they can represent complex relationships between input variables and power output. Malakouti

reviews the development of ML approaches in this field, including classical algorithms, deep-learning models, hybrid techniques, explainable AI, and emerging research directions [5]. Recurrent neural networks, including long short-term memory (LSTM) and gated recurrent unit (GRU) networks, are often considered for sequential data because they can learn temporal dependencies. Other architectures may be useful for capturing spatial relationships or more complex patterns [7], [8].

Recent work has also expanded from forecasting at a single site to prediction across multiple sites. Verdone et al. review solar- and wind-forecasting methods across both settings and discuss the importance of combining temporal information with data from several generation locations [6]. This multi-site perspective aligns with federated learning: participating sites can contribute to a shared forecasting model without sending their complete datasets to a central server.

Renewable-energy forecasting is consequently both a prediction task and a distributed-learning problem. A practical system must provide useful forecasts while accounting for data privacy, communication costs, available computing resources, and differences among individual sites. These considerations motivate the combination of ML/DL forecasting methods with federated and distributed learning.

C. Prosumer Analytics

Modern power systems increasingly include prosumers—users who both consume and produce electricity. A household with rooftop PV, battery storage, or an electric vehicle (EV), for example, may generate electricity, draw power from the grid, store energy, and adjust its demand. Similar roles are found in commercial buildings, industrial facilities, and community energy systems.

Prosumer analytics examines generation, consumption, storage, and operating data to understand energy use and support management decisions. Forecasts of prosumer generation and demand can help estimate future energy requirements, plan battery charging and discharging, schedule EV charging, and identify opportunities for demand response [2], [4].

These measurements can also be sensitive. Detailed electricity-use records may reveal information about the activities of people or the operation of a building. Centralizing such data can therefore create privacy and security concerns. Federated learning offers one possible approach: prosumers or their energy-management devices can train models locally and contribute updates to a shared model without directly sharing their raw measurements [2], [4].

The literature identifies potential federated-learning applications in both renewable generation and energy consumption. Grataloup et al. review work involving solar and wind energy, buildings, and other energy systems, highlighting data heterogeneity and privacy as important considerations [4]. Federated learning has also been studied for energy management and demand response, where participants can learn collaboratively while keeping their own data at their respective sites [2].

Combining prosumer analytics with renewable-energy forecasting could support a broader energy-management framework. Such a framework could forecast generation and consumption, examine battery and EV operation, identify flexible loads, and inform demand-response decisions.

D. Federated Learning Motivation

Federated learning has emerged in response to the growth of distributed datasets and the practical limits of centralized machine learning. In a centralized approach, data from multiple users or organizations are transferred to a common server. Federated learning instead allows each participant to retain its dataset locally while contributing to the training of a shared model [3].

A typical federated-learning round begins when a coordinator sends the current model to participating clients. Each client trains the model using its local data and returns an update. The coordinator aggregates the updates to produce a revised global model, which is then sent back to the clients. This process can continue over several communication rounds [3]. Since the raw datasets remain at the client sites, the approach reduces the need to transfer sensitive energy records.

In an energy application, clients could include PV plants, wind farms, smart buildings, households, battery systems, and EV charging stations. Their data are unlikely to have identical distributions: location, weather, generation capacity, equipment, and user behaviour all affect the measurements. This variation creates the non-independent and identically distributed (non-IID) data problem, a major challenge in federated learning [3], [4].

Other challenges reported in the renewable-energy literature include communication overhead, privacy protection, fairness, data heterogeneity, and scalability [4]. A standard federated-learning algorithm may not address all of these issues. Possible responses include personalized federated learning, client clustering, communication-efficient training, and additional privacy-preserving mechanisms.

This paper examines how distributed and federated learning can be combined with renewable-energy forecasting and prosumer analytics to support collaborative energy intelligence while keeping raw data at their source.

E. Scope

This paper focuses on the use of federated and distributed machine learning for renewable-energy forecasting and prosumer analytics. The discussion covers solar and wind forecasting, distributed energy-resource data, prosumer generation and consumption, battery and EV analytics, demand response, privacy, and federated model training.

The literature review considers eight selected studies on renewable-energy forecasting, ML and DL, distributed machine learning, federated learning, privacy-aware energy management, and renewable-energy applications. Drawing on these studies, the paper outlines an architecture in which distributed energy clients process their own data and train local models, while a federated coordinator aggregates model updates.

The paper also discusses practical issues such as non-IID data, communication overhead, privacy and security, personalization, scalability, and client reliability. The aim is to describe a framework for collaborative forecasting and prosumer analytics that does not require complete centralization of sensitive energy data.

II. LITERATURE REVIEW

A. Review of Eight Selected Papers

2.1 Review of Eight Selected Papers

Ref.	Focus	Contribution	Relevance
[1]	AI / smart grids	Advanced AI, forecasting and scenario analysis	Future intelligent forecasting
[2]	FL demand response	Privacy-preserving distributed management	Prosumer energy management
[3]	Distributed ML / FL	Architecture, aggregation, communication, security	FL foundation
[4]	FL renewable energy	Solar, wind, buildings and other applications	Direct topic relevance
[5]	ML renewable forecasting	Model families and future directions	Forecasting foundation
[6]	Solar/wind forecasting	Single-site to multi-site paradigm	Spatial/distributed forecasting
[7]	Deep-learning forecasting	Taxonomy of wind and solar DL	Model selection
[8]	Wind deep representation	Advanced wind representation learning	Wind forecasting

B. Machine Learning Forecasting

2.2 Machine Learning Forecasting

C. Distributed and Federated Learning

2.3 Distributed and Federated Learning

Federated learning allows multiple clients to develop a shared model without transferring

their raw observations. Each client optimizes the model using its own energy data and sends model parameters or updates to a coordinating server. The server aggregates the updates and distributes the revised global model to the clients. This arrangement is relevant to energy systems in which data are held by PV plants, wind farms, buildings, battery systems, EV charging stations, and household prosumers. A practical implementation must still address aggregation efficiency, communication, update security, parallel local training, and the possibility that some clients may be unavailable [3].

D. Privacy-Preserving Energy Management

2.4 Privacy-Preserving Energy Management

Demand-response applications show why local energy information may need to be protected. Consumption profiles can reveal sensitive details about users, buildings, and operating activities. Federated learning can support shared energy-management models while reducing the need to exchange these records directly [2]. Each participant trains a model using local consumption and operating data, then sends model updates to a coordinator. The aggregated model can provide shared information for demand prediction, load management, scheduling, and demand response while the original records remain at their respective sites. However, model updates may still reveal information about local training data, so additional privacy and security measures may be necessary [2].

E. Research Gap

2.5 Research Gap

Several issues remain open in federated learning for renewable-energy applications. These include non-IID data, communication overhead, possible privacy leakage, fairness, client reliability, scalability, and limited validation under real operating conditions. Renewable-energy clients differ in location, weather, generation capacity, consumption patterns, and available equipment. Such differences can make collaborative training more difficult. Repeated communication between clients and the coordinator can also increase bandwidth and training costs. Keeping raw data local does not, by itself, guarantee privacy, because information may still be inferred from model updates. Differences in computing resources and participation can further affect fairness and reliability. The reviewed literature therefore points to the need for robust, scalable, privacy-aware methods that can handle heterogeneous data, together with more field-based evaluation [4].

III. SYSTEM ARCHITECTURE

A. Overall Architecture

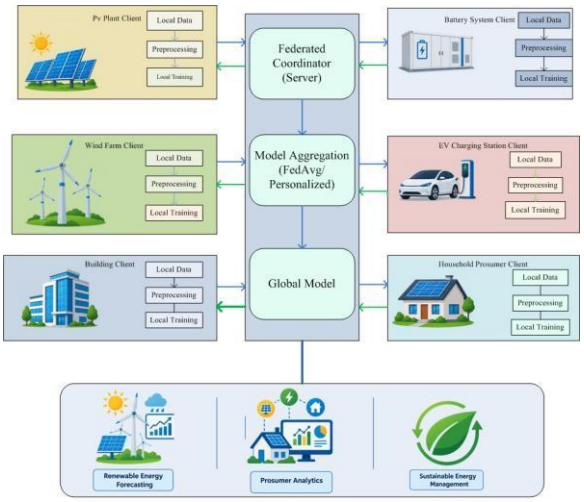

Fig. 1. Proposed federated and distributed architecture for renewable-energy forecasting and prosumer analytics.

3.1 Overall Architecture

Fig. 1 presents the proposed federated and distributed architecture for renewable-energy forecasting and prosumer analytics. The system includes several types of distributed clients: PV plants, wind farms, buildings, battery systems, EV charging stations, and household prosumers. Each client collects local energy and environmental measurements, preprocesses its data, and trains a model locally. Instead of sending raw records, clients transmit protected model updates to the federated coordinator. The coordinator aggregates the updates to form a global model and redistributes it for the next training round. In this way, the architecture supports collaborative learning while retaining local data ownership [3], [4].

B. Distributed Clients

The proposed clients may represent PV plants, wind farms, commercial buildings, battery systems, EV charging stations, and household prosumers. Their generation capacities, climate conditions, sensor configurations, load profiles, operating patterns, and network connectivity may differ. These differences allow the shared model to learn from a range of energy conditions, but they also create statistical heterogeneity, or non-IID data, because the feature and target distributions are not the same at every client [3], [4].

Measurements such as renewable generation, electricity consumption, weather variables, battery state of charge, and EV charging demand remain at the relevant site. Each client preprocesses its records and trains a forecasting model locally before sending model updates to

the coordinator. The clients can therefore contribute to a shared model without transferring their complete datasets [3], [4].

A single global model may not perform equally well for every client. Personalized federated learning can adapt the shared model to local conditions, while clustered federated learning can group clients with similar locations, operating conditions, or data distributions. Grataloup et al. discuss approaches such as local model components and location-based clustering for renewable-energy applications [4].

C. Local Data Pipeline

3.3 Local Data Pipeline

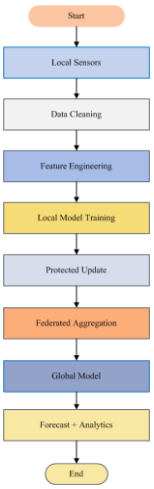

Fig. 2. Local data-to-federated forecasting flow.

Fig. 2 shows the path from locally collected energy measurements to federated forecasting. The process begins with energy and environmental data collection. Each client then handles missing values, noise, outliers, and time alignment, and applies normalization and time-window construction as required. Features may include lagged generation, weather measurements, rolling statistics, calendar information, and device states. The client uses these processed data to train its local forecasting model. It then sends protected model information to the federated coordinator rather than transferring the original dataset.

D. Federated Training

$w_k^* = \arg\min_v (1/n_k) \Sigma_{i=1}^{n_k} L(f_v(x_i), y_i)$ (1)

In this expression, w_k^* denotes the optimized local model, n_k is the number of samples held by client k, x_i is an input vector, y_i is the corresponding actual value, and L is the selected loss function.

$w_t h_1 = \Sigma_{k=1}^{K} (n_k/n)\ w_t h_1^k$ (2)

Federated training allows distributed energy clients to train a forecasting model together without sharing their raw datasets. During each communication round, every participating client optimizes the model using its local training samples, as represented by (1). After local training, the clients send their updated parameters to the coordinator. The coordinator

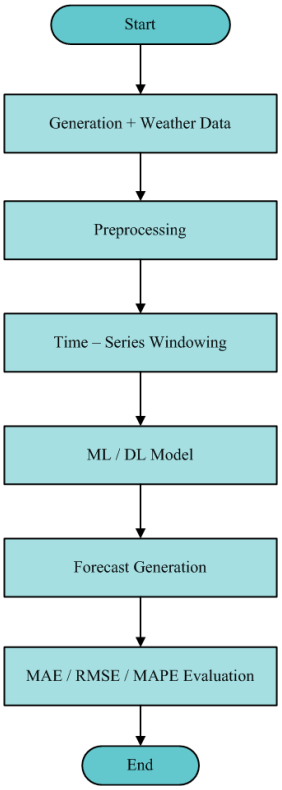

combines these updates using sample-weighted aggregation, as shown in (2), so clients with more training samples have a larger contribution to the updated global model. This enables data and computing resources from different sites to contribute to a common model while the original training data remain local [3]. However, differences in location, weather, generation, and consumption can produce non-IID data. Under these conditions, alternative aggregation methods, clustering, or personalization may be needed [4].

Fig. 3. Renewable-energy forecasting pipeline.

Fig. 3 outlines the renewable-energy forecasting process. Historical generation measurements and relevant weather or environmental variables are supplied to the trained forecasting model. The model learns temporal and nonlinear relationships in these inputs and estimates generation over a future time horizon. Depending on the application, candidate models may include LSTM, GRU, CNN-based, or Transformer architectures. The resulting forecasts can then inform energy-management decisions.

E. Forecasting and Prosumer Analytics

$\hat{y}_t h_h = f(X_t, W)$ (3)

The forecasting stage uses the trained federated model to estimate future renewable generation, electricity demand, or other energy-related quantities. Equation (3) expresses the general forecasting relationship: X_t represents the input features at time t, W denotes the trained model parameters, and $\hat{y}_{t+h}$ is the prediction for a future horizon h. Inputs may include historical generation, weather conditions, consumption patterns, and device states. Forecasts can support energy scheduling and the integration of renewable generation [4].

E_net = E_generation − E_consumption (4)

Prosumer analytics also examines the balance between locally generated and consumed energy. Equation (4) defines net energy as generation minus consumption. A positive value

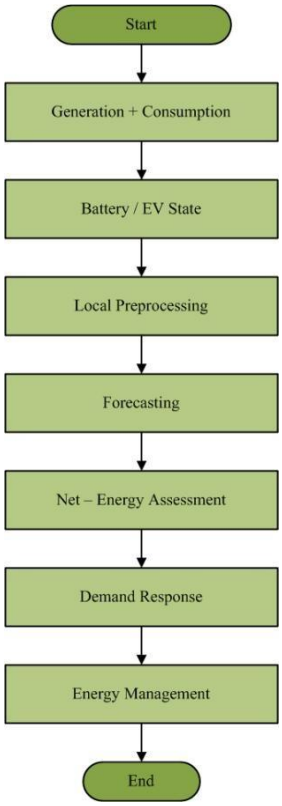

indicates a surplus that may be stored or exported, while a negative value indicates that demand exceeds local generation. This information can support battery scheduling, EV charging, demand response, and other energy-management decisions. Federated methods can enable such analyses across multiple prosumers without centralizing their sensitive records [4].

Fig. 4. Prosumer analytics flow.

Fig. 4 illustrates how forecasting outputs can be used in prosumer energy management. A prosumer may generate and consume electricity and may also operate battery storage or charge an EV. By considering renewable generation, consumption, battery status, and EV charging demand, the system estimates the prosumer's net-energy position. The resulting information can guide battery charging or discharging, EV charging schedules, renewable-energy use, and demand-response actions.

F. Security and Privacy

3.6 Security and Privacy

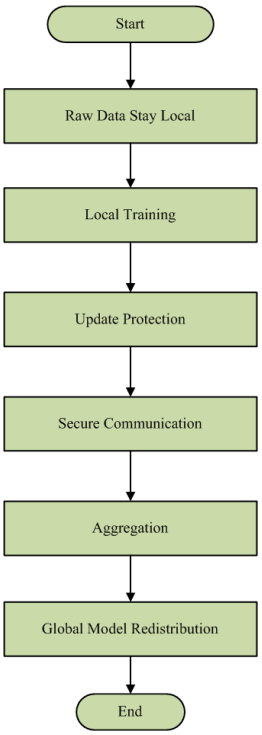

Fig. 5. Privacy-preserving federated learning flow.

Fig. 5 illustrates the privacy-aware learning process. Each client retains its raw energy data and trains a model locally. Before transmission, the resulting model information can be protected and sent through a secure communication channel to the federated coordinator. The coordinator aggregates updates from participating clients and distributes the revised global model. Depending on the threat model, protection mechanisms may include secure aggregation, encryption, authentication, and differential privacy. These measures complement data locality; they do not remove the need to assess the privacy risks associated with model updates.

IV. METHODOLOGY

A. Data Preparation

A future implementation should use multi-site renewable-energy datasets or construct realistic client partitions representing PV plants, wind farms, buildings, battery systems, EV charging stations, and household prosumers. Each client should contain time-series measurements such as renewable generation, electricity consumption, weather conditions, and device operating states. The data should be divided chronologically into training, validation, and test sets so that information from the future is not used during training. This is especially important for time-series forecasting, where random splitting can cause temporal leakage. Missing values, noise, and outliers should be treated before model training. Normalization and feature construction should preferably take place locally at each client, consistent with the data-locality principle of federated learning [3]. Client partitions should preserve realistic differences in location, weather, equipment, and operating conditions rather than making all datasets artificially identical [4].

B. Model Development

The model-development stage can compare a classical ML baseline with deep-learning and federated deep-learning models. Classical methods provide a useful reference, while LSTM, GRU, and Transformer architectures can be assessed for their ability to learn temporal dependencies and nonlinear patterns in renewable-energy data. LSTM and GRU models are commonly used for sequential forecasting, whereas Transformer models use attention mechanisms to capture longer-range dependencies [7], [8]. Federated versions should use the same client partitions and forecasting inputs as the comparison models. Model selection should consider not only prediction accuracy but also training time, communication requirements, memory use, and suitability for edge devices [5]. Evaluation can use MAE, RMSE, and MAPE alongside computational and deployment measures.

Simulation and modelling can be used to examine different renewable-energy integration scenarios. For example, grid configurations with varying levels of renewable penetration can be studied to assess voltage stability, power quality, and system reliability. Models may also include energy storage and smart-grid technologies intended to manage the variability of renewable generation. Such simulations can help examine system behaviour and compare possible operating strategies. The specific models, assumptions, and simulation results should be reported when an implementation is carried out.

$$MAE = (1/N) \Sigma_{i=1}^{N} |y_i - \hat{y}_i| \quad (5)$$

$$RMSE = \sqrt{[(1/N) \Sigma_{i=1}^{N} (y_i - \hat{y}_i)^2]} \quad (6)$$

$$MAPE = (100/N) \Sigma_{i=1}^{N} |(y_i - \hat{y}_i)/y_i| \quad (7)$$

MAE measures the average absolute difference between predicted and actual values. RMSE gives greater weight to larger errors, while MAPE expresses the error as a percentage. MAPE can become unstable when actual values are zero or close to zero. The evaluation metric should therefore be selected with the characteristics of the renewable-energy dataset in mind.

C. Federated Optimization

The study should compare local-only, centralized, standard federated, and personalized federated training to examine how the learning strategy affects renewable-energy forecasting. A local-only model provides a baseline in which each client learns independently. Centralized training can be included when data-sharing restrictions allow it. Standard federated learning supports joint training without transferring local datasets, while personalized federated learning adapts shared knowledge to individual clients [3], [4].

The effect of non-IID data can be studied by partitioning clients according to geographic location, renewable-generation capacity, weather conditions, consumption behaviour, or equipment type. These differences change the statistical distributions of the local datasets and may reduce the performance of conventional FedAvg [4]. Personalization layers, client clustering, and local fine-tuning can be evaluated as alternatives to relying on one global model.

D. Experimental Plan

A fair experimental comparison should use the same datasets, forecast horizons, preprocessing steps, and test periods for local, centralized, federated, and personalized federated models. This helps ensure that observed differences arise from the learning strategy rather than inconsistent experimental conditions. Results should be reported for individual clients as well as for the overall average, since heterogeneous sites may experience different forecasting performance [3], [4]. Evaluation should include MAE, RMSE, and MAPE, together with communication rounds and computational requirements where possible. Repeating experiments with multiple random seeds can help assess stability and reproducibility. Sensitivity tests can examine the effects of client count, local training epochs, aggregation strategy, forecast horizon, and degree of data heterogeneity. This design is consistent with studies that compare local, centralized, federated, and personalized approaches across different forecasting horizons [4].

A later stage of the work could use small-scale pilot implementations to assess practical feasibility. For example, a pilot could combine renewable generation, energy storage, and

smart-grid components under selected operating conditions. Measurements from these trials would help evaluate performance, identify implementation challenges, and determine whether the approach can be scaled. These activities are proposed as future validation; they should not be interpreted as experiments already completed in this paper.

V. RESULTS AND DISCUSSION

A. Literature-Based Results

The proposed framework brings together federated and distributed learning for renewable-energy forecasting and prosumer analytics while retaining data at the participating sites. The reviewed literature reports that ML and DL models can capture nonlinear and temporal patterns in solar and wind generation data [5], [7], [8]. Federated learning provides a way for multiple energy clients to contribute to a shared model without directly transferring their raw datasets [3], [4]. A future benchmark could compare local-only, centralized, standard federated, and personalized federated models using consistent forecasting horizons and test periods. Relevant measures include MAE, RMSE, MAPE, communication cost, and computational requirements. Client-level evaluation is important because sites differ in weather, generation capacity, sensors, and consumption patterns. Under non-IID conditions, personalization or clustering may help adapt shared learning to individual clients [4]. The reviewed sources do not establish one numerical accuracy value for the proposed architecture; its performance must be measured through a future implementation using consistent datasets and evaluation procedures.

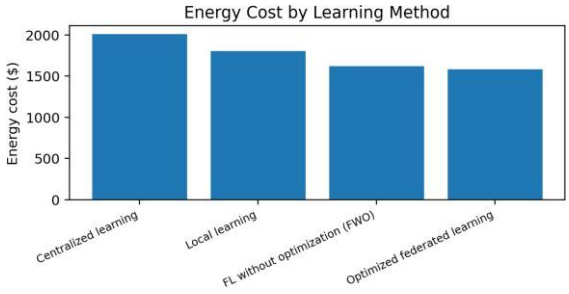

Fig. 6. Energy cost by learning method, based on literature-reported values from the reviewed demand-response study [2].

The comparison in Fig. 6 is based on energy-cost values reported in the reviewed demand-response study. It indicates a lower reported energy cost for optimized federated learning than for centralized learning. These values are taken from the cited study and are not results from a new experiment conducted in this paper [2].

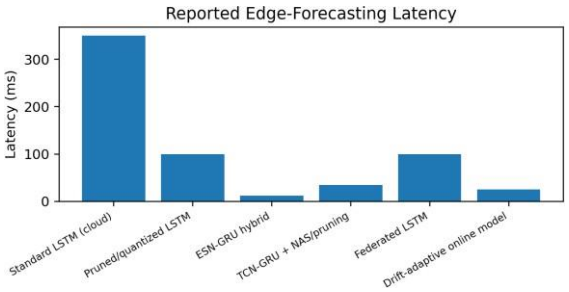

Fig. 7. Reported edge-forecasting latency for selected models, compiled from the forecasting review [5].

Fig. 7 compares edge-forecasting latency values reported in the forecasting review. The comparison suggests that model compression and lightweight architectures can reduce latency relative to a standard cloud-based LSTM. These literature-derived values illustrate the trade-off between forecasting capability and the resource limits of edge devices [5].

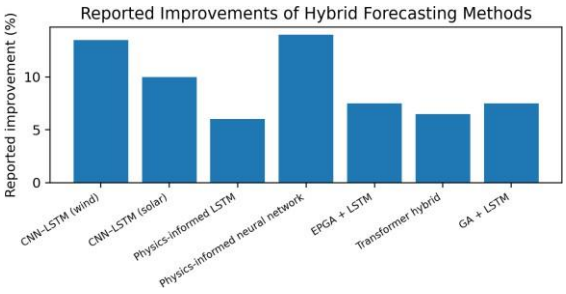

Fig. 8. Reported improvements of selected hybrid forecasting methods; underlying metrics differ [5].

Fig. 8 summarizes percentage improvements reported for selected hybrid forecasting methods. The studies use different evaluation measures, including RMSE, MAE, nRMSE, NSE, and MAPE. The plotted values should therefore be read as results reported within individual studies, not as a direct comparison under one common benchmark [5].

B. Discussion

The reviewed studies suggest that federated learning could provide a useful framework for combining distributed renewable-energy forecasting with prosumer analytics. One important benefit is that participating clients can keep their raw energy records locally while contributing to collaborative model training [3], [4]. This is relevant when measurements contain sensitive operational or consumption information. However, federated learning does not automatically eliminate privacy risks: model updates may still reveal information about local training data, so suitable protection mechanisms remain necessary [4].

Data heterogeneity is another important consideration. Solar plants, wind farms, buildings, and households operate under different environmental and behavioural conditions, producing non-IID datasets. A single global model may therefore perform unevenly across clients. Personalized federated learning, clustering, and local adaptation are possible ways to address these differences [4]. Model choice also involves a trade-off between forecasting performance and computational cost. LSTM, GRU, and Transformer models offer different capabilities for temporal prediction, but their suitability depends on the resources available at each edge device [5], [7], [8]. The proposed architecture should be tested with consistent datasets, forecast horizons, client partitions, and test periods before quantitative claims are made. The literature supports the motivation for the framework, but experiments are still needed to measure its accuracy, communication efficiency, scalability, and personalization benefits.

VI. CONCLUSION

Federated and distributed machine learning offer a possible foundation for privacy-aware renewable-energy forecasting and energy management. PV plants, wind farms, battery systems, EV charging stations, and household prosumers generate data that vary with location, weather, consumption patterns, and equipment characteristics. Collecting all of these records at a central server can create privacy, security, communication, and scalability concerns. Federated learning addresses part of this challenge by allowing clients to train models locally and send model updates to a coordinator. The coordinator aggregates those updates to produce a shared model without directly receiving the raw datasets.

Distributed learning also allows geographically separated devices and energy resources to contribute to a common learning process. In combination with deep-learning models such as LSTM, GRU, and Transformer networks, these approaches can be applied to forecasting solar and wind generation, electricity demand, battery behaviour, and prosumer energy exchange. The reviewed literature indicates that this direction is promising for renewable-energy integration and intelligent energy management. However, the proposed architecture still requires implementation and evaluation using suitable datasets and consistent experimental procedures. Future work should quantify forecasting performance, communication costs, privacy protection, and the benefits of personalization under realistic operating conditions.

REFERENCES

1. “Generative AI and LLM applications in renewable energy and smart grids,” systematic review paper, 2025.
2. M. Alhazmi, A. P. F. Zhao, W. Li, and C. Yang, “Federated Learning for Real-Time Demand Response by Data Centers Toward Energy Efficiency and Privacy Preservation,” IEEE Access,

2025, doi: 10.1109/ACCESS.2025.3583436.

3. J. Liu, J. Huang, Y. Zhou, X. Li, S. Ji, H. Xiong, and D. Dou, “From Distributed Machine Learning to Federated Learning: A Survey,” Knowledge and Information Systems, 2022.
4. Grataloup, S. Jonas, and A. Meyer, “A review of federated learning in renewable energy applications: Potential, challenges, and future directions,” Energy and AI, vol. 17, 2024, Art. no. 100375.
5. M. Malakouti, “Machine learning for renewable energy forecasting: advances, challenges, and future directions,” Sustainable Energy Research, vol. 13, 2026, Art. no. 14, doi: 10.1186/s40807-026-00242-x.
6. Verdone et al., “A review of solar and wind energy forecasting: From single-site to multi-site paradigm,” Applied Energy, vol. 392, 2025, Art. no. 126016.
7. “A review and taxonomy of wind and solar energy forecasting methods based on deep learning,” Energy and AI, vol. 4, 2021, Art. no. 100060.
8. “Wind power forecasting with deep representation: Asurvey,” Applied Energy,2026 doi:10.1016/j.apenergy.2026.128106.