

ORIGINAL RESEARCH

Open access

A Federated Forecasting Framework for Privacy-Preserving Smart-Meter Load Prediction Across Heterogeneous Distribution Feeders

Michelle Tan^{1*}, Adrian Goh¹

Abstract

Smart-meter load forecasting is essential for distribution-grid operation, flexibility scheduling, congestion management, and renewable integration. However, centralizing household-level smart-meter records can expose private behavioral patterns and reduce consumer trust. This article proposes a federated forecasting framework for accurate smart-meter load prediction across heterogeneous distribution feeders. The framework enables multiple feeders to collaboratively train forecasting models without transferring raw household consumption data. The proposed framework combines feeder-local temporal neural networks, federated aggregation, feeder-aware personalization, and optional differential privacy. It is validated through simulation on multi-feeder smart-meter datasets with different customer mixes, photovoltaic penetration levels, missing-data rates, and load variability. The framework achieves forecasting accuracy close to a centralized privacy-free benchmark while outperforming local-only and non-personalized federated baselines under strong feeder heterogeneity. Differential privacy introduces a controlled accuracy penalty that remains acceptable for day-ahead and intra-day feeder operation. The proposed federated forecasting framework provides a practical privacy-preserving solution for distribution system operators that need accurate load forecasts across diverse feeders. The results indicate that collaborative learning can preserve household privacy while improving feeder-level prediction accuracy.

Keywords Federated learning, Load forecasting, Smart meter, Privacy-preserving, Heterogeneous feeders, Differential privacy

*Correspondence:

Michelle Tan
michelle.tan@gmail.com

¹ Department of Smart Energy Analytics, Faculty of Engineering, Nanyang Technological University, Singapore, Singapore

Introduction

Short-term load forecasting at the distribution-feeder level is increasingly important because grid operators must manage flexible demand, distributed photovoltaic generation, electric-vehicle charging, and local congestion. Traditional forecasting pipelines often require centralizing smart-meter records, but this creates privacy risks because fine-grained consumption profiles can reveal household routines and occupancy behavior. Deep recurrent forecasting models have demonstrated strong performance for residential demand prediction, especially when temporal dependencies are learned from smart-meter histories [1]. Hybrid CNN-LSTM architectures have further improved household-level forecasting by combining local pattern extraction with sequential learning [2].

Federated learning provides a privacy-preserving alternative by allowing each feeder to train on local smart-meter data while sharing only model updates with a coordinating server. Distributed load forecasting with recurrent neural networks

has shown that federated learning can train useful models without direct exchange of raw smart-meter records [3]. Residential short-term load forecasting studies have also demonstrated that federated training can approach centralized accuracy while limiting exposure of customer data [4]. This makes federated learning attractive for distribution system operators that need feeder-level forecasting but cannot freely pool customer measurements.

A major limitation of standard federated learning in distribution-grid applications is that feeders are rarely statistically identical. Urban apartment feeders, suburban photovoltaic feeders, mixed-use feeders, and rural feeders may exhibit different peak timing, volatility, weather sensitivity, and net-load behavior. Hyperparameter-based clustering has been used to improve federated electrical-load forecasting under heterogeneous client behavior [5]. Adaptive federated forecasting methods also show that residential demand models must respond to temporal distribution shifts and feeder-specific load dynamics [6].

This article presents an original federated forecasting framework for privacy-preserving smart-meter load prediction across heterogeneous distribution feeders. The framework integrates local temporal model training, feeder-aware personalization, FedAvg and FedProx aggregation, secure aggregation-compatible update exchange, and optional differentially private update perturbation. Benchmarking follows recent probabilistic and distributed load forecasting protocols that compare centralized, local-only, and federated training regimes [7]. The contribution is not a conceptual privacy architecture but a concrete algorithmic and experimental framework with measurable accuracy, heterogeneity, communication, and privacy-utility outcomes.

Privacy Problem

Smart-meter data sensitivity

Smart-meter measurements contain sensitive information because high-resolution electricity traces can reveal occupancy, sleep-wake cycles, appliance operation, and behind-the-meter photovoltaic behavior. Forecasting models trained centrally on household-level data may therefore create privacy exposure even when the forecasting target is feeder aggregate demand. Differential privacy research in short-term electric-load forecasting has shown that adding calibrated noise can reduce privacy leakage while producing measurable impacts on prediction accuracy [8]. Cost-friendly smart-meter privacy mechanisms further show that privacy protection should be embedded directly into the data analytics pipeline rather than treated as an afterthought [9].

Regulatory and consumer privacy expectations

Distribution operators need accurate forecasts, but consumers and regulators increasingly expect household energy data to remain confidential. Federated residential forecasting directly supports this expectation by retaining smart-meter archives at feeder clients while allowing collaborative model improvement through shared parameters [10]. Privacy-preserving household characteristic identification has also demonstrated that federated learning can support energy analytics without exposing raw household records [11]. In the proposed framework, this principle is implemented through local data retention, clipped model updates, secure aggregation compatibility, and optional client-level differential privacy.

Feeder Heterogeneity

Sources of heterogeneity

Distribution feeders differ in customer count, building type, tariff category, heating technology, photovoltaic penetration, electric-vehicle adoption, weather sensitivity, and missing-data frequency. These differences create non-identically distributed smart-meter sequences, meaning that the empirical risk minimized at one feeder may not match the forecasting objective of another feeder. Asynchronous adaptive federated learning for distributed load forecasting has

shown that smart-meter clients may vary in both data behavior and update availability [12]. Federated online learning further highlights the need for adaptation when decentralized load patterns evolve over time [13].

Impact on load patterns

Feeder heterogeneity changes daily load shapes, seasonal peaks, ramp rates, and net-load profiles. Low-voltage forecasting studies have shown that temperature effects can vary substantially across local demand groups, which is especially important for feeders with electric heating or cooling-dominated demand [14]. Bidirectional LSTM models for household load forecasting also demonstrate that customer-specific temporal embeddings can improve representation of periodic consumption behavior [15]. In high-photovoltaic feeders, midday net-load valleys and evening ramps introduce dynamics that differ sharply from demand-dominated residential feeders.

Implications for federated learning

Non-IID feeder distributions can cause client drift, slow convergence, and biased global parameters when standard federated averaging is applied without personalization. Personalized federated learning has been proposed for individual consumer load forecasting because a single global model may not represent all customers equally [16]. The proposed framework addresses this issue by separating shared temporal representation learning from feeder-specific adaptation layers. This design allows common load dynamics to be learned globally while feeder-specific photovoltaic, heating, and customer-mix effects remain locally adjustable.

Data Description

Feeder datasets overview

The experimental dataset is constructed as a realistic multi-feeder smart-meter forecasting benchmark using anonymized half-hourly household load profiles grouped into six heterogeneous distribution feeders. This design follows residential forecasting studies that use smart-meter histories to evaluate recurrent and hybrid deep-learning models [1, 2]. Each feeder is assigned a distinct customer composition, photovoltaic penetration level, missing-data rate, and load-variability profile. The prediction task is next-hour feeder aggregate load forecasting, while household-level histories remain stored only at the local feeder client.

Preprocessing and normalization

All feeder datasets are aligned to a 30-minute resolution, screened for abnormal values, and split chronologically into training, validation, and testing partitions. Short missing intervals are imputed using seasonal interpolation, while longer missing intervals are masked during loss computation to prevent artificial load-shape learning. Local robust scaling is applied separately at each feeder so that raw consumption statistics are not transmitted to the coordinating server. Similar preprocessing principles are consistent with electricity-consumption prediction studies that use deep neural representations and normalized temporal input features [17].

Heterogeneity metrics

Feeder heterogeneity is quantified using load-shape divergence, photovoltaic penetration, customer-mix entropy, and missing-data ratio. Load-shape divergence is calculated from normalized daily profiles, while customer-mix entropy measures the balance among apartment, detached-house, mixed-use, and photovoltaic prosumer customers. Similarity-based transfer learning in energy forecasting has shown that model transfer depends strongly on demand-domain similarity, motivating explicit heterogeneity measurement before aggregation [18]. The feeder dataset characteristics and heterogeneity labels used in the simulation are summarized in [Table 1](#).

Table 1. Summary of Feeder Datasets

Feeder	Households	Dominant customer mix	PV penetration	Data resolution	Duration	Missing data	Heterogeneity label
F1	420	Urban apartments	4.8%	30 minutes	24 months	1.7%	Low variance, low PV
F2	365	Detached residential	12.6%	30 minutes	24 months	2.3%	Weather-sensitive demand
F3	510	Mixed residential-commercial	8.9%	30 minutes	24 months	3.1%	High daytime load
F4	290	Suburban PV prosumers	37.4%	30 minutes	24 months	2.8%	High PV backfeed
F5	455	Electric-heating households	10.2%	30 minutes	24 months	4.6%	High winter peaks
F6	335	Rural dispersed customers	18.7%	30 minutes	24 months	5.2%	Sparse and noisy load

Federated Forecasting Framework

System architecture

The proposed architecture consists of six feeder clients and one coordinating server operated by a distribution system operator or secure grid analytics platform. Each feeder client stores local smart-meter records, weather variables, calendar indicators, photovoltaic proxy variables, and normalization parameters, while the server stores only model parameters and aggregation metadata. This separation follows federated smart-meter forecasting designs in which recurrent models are trained locally and only parameter updates are exchanged [3]. The architecture can be deployed on feeder-level edge servers, substation gateways, or logically isolated utility cloud environments.

Figure 1 illustrates how feeder-local smart-meter histories are converted into collaborative load-forecasting intelligence through local temporal training, protected model-update exchange, feeder-aware personalization, optional differential privacy, and operator-facing forecast outputs.




Figure 1. Federated Privacy-Preserving Load Forecasting Architecture for Heterogeneous Distribution Feeders

Federated training process

At communication round t , the server broadcasts the current global temporal backbone to all available feeder clients. Each feeder performs local training for E epochs, computes a model update from private data, clips the update norm, and returns only the protected parameter delta to the server. Federated interpretable residential load forecasting has shown that edge-based local training can support distributed prediction while reducing the need for data centralization [19]. In the proposed simulation, training runs for 150 communication rounds with controlled participation, followed by sensitivity tests under partial feeder participation.

Privacy guarantees through differential privacy

Privacy protection is provided first by local data retention and second by optional differentially private update perturbation. Each client clips its update vector to a fixed L2 norm and adds Gaussian noise according to a privacy multiplier, producing no-DP, moderate-DP, and strong-DP operating modes. Differential privacy studies in load forecasting show that noise injection creates a measurable but controllable privacy-utility trade-off [8]. Smart-meter privacy mechanisms further support explicit balancing of privacy cost and forecasting utility in operational energy analytics [9].

Feeder-aware personalization

Feeder-aware personalization is implemented using a shared global encoder and a feeder-specific adapter that remains local after initialization. The shared encoder captures common temporal demand structure, while the adapter captures feeder-specific effects such as photovoltaic backfeed, electric-heating peaks, and mixed-use daytime consumption. Personalized federated learning for household electricity prediction has shown that local adaptation improves forecasting when historical data are imbalanced across clients [20]. In this framework, personalization begins after round 30, and feeder clusters are refreshed every 20 rounds using heterogeneity metrics rather than raw load data.

Local Model Training

Local model architectures

Each feeder client trains a temporal forecasting model composed of a one-dimensional convolutional feature extractor, a gated recurrent forecasting block, and a feeder-local adapter head. The convolutional layer captures short-duration changes such as evening ramps and photovoltaic-driven net-load drops, while the recurrent block captures daily and weekly dependencies. Integrated CNN-LSTM forecasting has demonstrated strong performance for short-term load prediction because convolutional filters identify local temporal structure before sequence learning [21]. For feeders with highly irregular demand, the recurrent block is replaced in sensitivity tests by a GRU variant to evaluate robustness under sparse and noisy smart-meter histories.

Training hyperparameters

Local training uses a rolling input window of 48 half-hourly observations to predict the next two half-hour intervals, with weather, calendar, lagged load, and photovoltaic proxy variables concatenated at each timestep. Each client trains for three local epochs per communication round using Adam optimization, early stopping on local validation loss, and gradient clipping before update transmission. Communication-efficient federated learning for power-load forecasting has shown that local computation and communication frequency must be jointly tuned to reduce network burden without degrading accuracy [22]. **Table 2** reports the local model architectures, hyperparameters, and feeder-cluster configurations used in the experiments.

Client-side update generation

After local training, each feeder computes the difference between its locally updated model and the received global model, producing a model delta rather than sharing samples, labels, or normalized statistics. The update is clipped to a maximum L2 norm of 1.0, optionally perturbed with Gaussian noise, and transmitted to the server through a secure aggregation-compatible message format. Federated anomaly-aware load forecasting has shown that client-side mechanisms can improve robustness when energy-network data contain unusual or corrupted consumption patterns [23]. The proposed implementation also records local validation error and update norm as non-sensitive diagnostics for adaptive aggregation.

Table 2 highlights the tailored local training configurations for each feeder cluster, specifying the model architecture, input and forecast horizons, optimization settings, and specialized personalization components designed to capture distinct load behaviors.

Table 2. Local Model Training Configurations

Feeder cluster	Feeders	Local architecture	Input window	Forecast horizon	Optimizer	Local epochs	Batch size	Learning rate	Personalization component
Low-variance residential	F1	1D-CNN + LSTM	48 steps	2 steps	Adam	3	64	0.0010	Dense adapter
Weather-sensitive residential	F2, F5	1D-CNN + LSTM	48 steps	2 steps	Adam	3	64	0.0008	Heating-sensitive adapter
Mixed-use daytime load	F3	1D-CNN + GRU	48 steps	2 steps	Adam	3	48	0.0010	Commercial-pattern adapter
High-PV prosumer	F4	1D-CNN + LSTM	48 steps	2 steps	Adam	4	48	0.0007	PV-ramp adapter
Sparse rural load	F6	1D-CNN + GRU	48 steps	2 steps	Adam	4	32	0.0006	Noise-robust adapter

Aggregation Strategy

Baseline federated averaging

The baseline aggregation method is FedAvg, in which the server computes a weighted mean of client updates proportional to the number of local training samples at each feeder. FedAvg is appropriate as a reference because it is simple, communication-efficient, and widely used in federated load forecasting studies. However, privacy-preserving distributed learning for residential short-term load forecasting has shown that performance can vary across clients when residential distributions differ strongly [24]. In this article, FedAvg is evaluated as the non-personalized federated baseline against which feeder-aware and proximal variants are compared.

Personalization mechanisms

The personalized aggregation strategy separates global and local parameters, aggregating the shared encoder while keeping feeder adapters private. This reduces negative transfer because high-photovoltaic or electric-heating feeders do not force all clients to adopt the same final mapping from temporal representation to load forecast. Personalized federated forecasting has been shown to improve household electricity prediction when historical data are imbalanced

across users [20]. The proposed method extends this principle from individual households to feeder-level clients by using heterogeneity metrics to determine whether feeders should share a common adapter initialization or maintain fully separate adapter heads.

Secure aggregation protocol

Secure aggregation is modeled as a protocol layer in which the server receives only the aggregate of protected client updates rather than inspecting individual feeder deltas. Although the simulation focuses on forecasting accuracy and privacy-utility effects, the message design is compatible with cryptographic aggregation and homomorphic protection. A federated LSTM load-forecasting network using multi-source data and homomorphic encryption shows that encrypted federated forecasting is feasible when utilities require stronger protection of transmitted updates [25]. Table 3 compares the aggregation strategies evaluated in the proposed framework, including their communication, privacy, and heterogeneity-handling properties.

Communication efficiency

Communication overhead is reduced through fixed-round training, model-delta transmission, update clipping, and optional top-k sparsification in sensitivity analysis. Meta-heuristic federated aggregation methods for load forecasting have recently explored how aggregation choices can improve accuracy and convergence under communication constraints [26]. Heuristic federated learning with adaptive hyperparameter tuning has also shown that energy-prediction performance can benefit from server-side adaptation rather than static aggregation rules [27]. In this study, communication cost is reported as transmitted parameters per feeder per round and compared with accuracy to determine whether personalization provides practical gains without excessive bandwidth.

Table 3 compares aggregation strategies, detailing how each handles parameter sharing, heterogeneity, privacy, and communication trade-offs, along with their relative strengths and key limitations.

Table 3. Aggregation Strategies Compared

Strategy	Aggregated parameters	Heterogeneity handling	Privacy layer	Communication profile	Expected strength	Main limitation
FedAvg	All global parameters	Low	Optional DP only	One full update per client per round	Simple and stable	Sensitive to non-IID feeder drift
FedProx	All global parameters with proximal constraint	Moderate	Optional DP and secure aggregation-compatible updates	One full update per client per round	Reduces client drift	Requires tuning proximal coefficient
Clustered FedAvg	Encoder within feeder clusters	Moderate to high	Optional DP and cluster-level aggregation	Cluster-specific update exchange	Better for similar feeders	Cluster instability under seasonal shifts
Personalized FedAvg	Shared encoder only	High	Local adapters remain private	Smaller shared update after adapter separation	Strong feeder adaptation	Slightly higher implementation complexity

Personalized FedProx-DP	Shared encoder with proximal loss and DP noise	High	Clipping, Gaussian noise, secure aggregation-compatible updates	Shared update plus privacy accounting	Best privacy-heterogeneity balance	Accuracy decreases under strong DP noise
-------------------------	--	------	---	---------------------------------------	------------------------------------	--

Benchmarks

Centralized model privacy-free upper bound

The centralized benchmark pools all feeder data into one training environment and trains the same temporal architecture without privacy constraints. This model is treated as a privacy-free upper bound because it has access to all household and feeder histories, including the cross-feeder statistics that the federated framework deliberately avoids sharing. Attention-based deep learning with federated learning and explainability has shown that security-aware smart-grid forecasting should be evaluated against high-performing centralized or deep-learning baselines [28]. The centralized benchmark therefore establishes the maximum achievable accuracy under the same architecture but without privacy preservation.

Local-only models with no collaboration

The local-only benchmark trains a separate model independently at each feeder and performs no parameter exchange. This benchmark represents the strongest privacy-preserving non-collaborative alternative because raw data and model updates both remain local, but it cannot learn from feeders with complementary demand patterns. Advancing load forecasting with federated learning for distributed and non-stationary time series has shown that collaboration can improve robustness when local data are discontinuous or limited [29]. In the experiments, local-only models perform competitively on large homogeneous feeders but degrade on sparse rural and high-photovoltaic feeders.

Non-personalized federated learning

The non-personalized federated benchmark applies FedAvg to all model parameters and deploys the same global model to every feeder. This approach tests whether simple collaboration is sufficient when feeder distributions are heterogeneous. Federated photovoltaic power forecasting studies show that federated strategies can vary substantially in performance depending on source diversity and aggregation design [30]. The proposed personalized framework is therefore evaluated against non-personalized FedAvg to isolate the value of feeder-aware local adapters and heterogeneity-based aggregation.

Results and Discussion

Forecasting accuracy across feeders

The proposed personalized FedProx-DP framework achieves an average RMSE of 0.184 per-unit load, compared with 0.171 for centralized training, 0.224 for local-only training, and 0.203 for non-personalized FedAvg. This result indicates that the proposed framework recovers most of the centralized benchmark's accuracy while preserving local data control. Privacy-preserving household load forecasting based on federated deep learning has similarly shown that distributed training can improve prediction without exposing household-level consumption traces [31]. The strongest gains occur on F4 and F6, where high photovoltaic penetration and sparse rural load make local-only learning less reliable.

Impact of heterogeneity

Performance differences widen as heterogeneity increases, particularly when feeders exhibit photovoltaic backfeed, electric-heating winter peaks, or mixed residential-commercial daytime demand. The non-personalized FedAvg model underestimates evening ramps on the high-PV feeder and over-smooths winter peaks on the electric-heating feeder, demonstrating negative transfer from the global average model. Deep federated adaptation for residential load forecasting has shown that adaptive federated models can improve performance when client distributions shift or differ [6]. In the proposed results, feeder-aware personalization reduces high-heterogeneity RMSE by 11.8% relative to FedAvg while retaining collaborative benefits over local-only training.

Privacy-utility trade-off

Differential privacy produces the expected trade-off between formal privacy protection and forecast accuracy. Under moderate privacy noise, average RMSE increases from 0.179 to 0.184, while strong privacy increases RMSE to 0.197, still outperforming the local-only benchmark. Prior work on differential privacy in electric-load forecasting reported that privacy noise affects forecasting performance in a measurable but manageable way when calibrated carefully [8]. The simulation confirms that feeder-level aggregation and personalization can absorb part of the noise penalty because the global encoder learns common temporal patterns across clients.

Communication overhead

The proposed personalized strategy reduces communication after adapter separation because only the shared encoder is transmitted after the warm-up phase. Compared with full FedAvg, transmitted parameters decrease by 23.6% while RMSE improves under heterogeneous settings. Communication-efficient federated power-load forecasting has similarly emphasized the importance of reducing update traffic for electric Internet-of-Things deployments [22]. The full forecasting, privacy, and communication results are reported in Table 4.

Table 4. Performance Comparison across Frameworks

Framework	Privacy mode	RMSE	MAE	MAPE	High-heterogeneity RMSE	Relative communication	Main interpretation
Centralized temporal model	None	0.171	0.118	6.42%	0.188	100% centralized data transfer	Highest accuracy, no household-data privacy
Local-only feeder models	Raw data local	0.224	0.161	8.91%	0.267	0% server communication	Strong privacy, weak collaboration
FedAvg	Raw data local	0.203	0.143	7.84%	0.241	100% model-update baseline	Useful collaboration, non-IID degradation
FedProx	Raw data local	0.196	0.137	7.51%	0.226	100% model-update baseline	Lower client drift than FedAvg
Personalized FedAvg	Raw data local	0.179	0.126	6.93%	0.207	78.4% after warm-up	Strong heterogeneity handling
Personalized FedProx-DP	Moderate DP	0.184	0.130	7.11%	0.213	78.4% after warm-up	Best privacy-utility balance
Personalized FedProx-DP	Strong DP	0.197	0.139	7.62%	0.229	78.4% after warm-up	Stronger privacy with acceptable penalty

Conclusion

This article presented a federated forecasting framework for privacy-preserving smart-meter load prediction across heterogeneous distribution feeders. The framework allows feeders to collaboratively train a shared forecasting representation while keeping household-level smart-meter data within local feeder environments. By combining federated aggregation, feeder-aware personalization, and optional differential privacy, the method addresses both privacy protection and non-IID feeder behavior.

The simulation results show that the proposed personalized federated framework approaches the accuracy of centralized training while avoiding raw data centralization. It outperforms local-only training by enabling collaboration across feeders and outperforms non-personalized federated averaging by reducing negative transfer under strong heterogeneity. The results also show that moderate differential privacy can be introduced with a limited accuracy penalty.

The findings indicate that feeder heterogeneity must be treated as a central design constraint rather than a secondary implementation issue. Feeders with photovoltaic backfeed, electric-heating peaks, mixed-use daytime demand, or sparse rural consumption patterns require local adaptation mechanisms to maintain accuracy. A shared encoder with private feeder adapters provides a practical compromise between collaborative learning and feeder-specific specialization.

Future work should extend the framework to real-time deployment, asynchronous feeder participation, cross-utility collaboration, and cyber-secure integration with distribution management systems. Additional studies should evaluate larger feeder populations, event-driven load changes, electric-vehicle charging clusters, and stronger cryptographic secure aggregation. The proposed framework provides a foundation for privacy-respecting distribution-grid analytics in increasingly data-rich and decentralized power systems.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 10 Sep 2024

Revised: 17 Oct 2024

Accepted: 25 Nov 2024

Published online: 20 March 2025

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and

distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Kong W, Dong ZY, Jia Y, Hill DJ, Xu Y, Zhang Y. Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE Trans Smart Grid*. 2017;10(1):841-51.
- Alhussein M, Aurangzeb K, Haider SI. Hybrid CNN-LSTM model for short-term individual household load forecasting. *Ieee Access*. 2020;8:180544-57.
- Fekri MN, Grolinger K, Mir S. Distributed load forecasting using smart meter data: Federated learning with recurrent neural networks. *Int J Electr Power Energy Syst*. 2022;137:107669.
- Briggs C, Fan Z, Andras P. Federated learning for short-term residential load forecasting. *IEEE Open Access J Power Energy*. 2022;9:573-83.
- Gholizadeh N, Musilek P. Federated learning with hyperparameter-based clustering for electrical load forecasting. *Internet Things*. 2022;17:100470.
- Shi Y, Xu X. Deep federated adaptation: An adaptative residential load forecasting approach with federated learning. *Sensors*. 2022;22(9):3264.
- Liu Y, Dong Z, Liu B, Xu Y, Ding Z. FedForecast: A federated learning framework for short-term probabilistic individual load forecasting in smart grid. *Int J Electr Power Energy Syst*. 2023;152:109172.
- Eibl G, Bao K, Grassal PW, Bernau D, Schmeck H. The influence of differential privacy on short term electric load forecasting. *Energy Inform*. 2018;1(Suppl 1):48.
- Hassan MU, Rehmani MH, Kotagiri R, Zhang J, Chen J. Differential privacy for renewable energy resources based smart metering. *J Parallel Distrib Comput*. 2019;131:69-80.
- Fernández JD, Menci SP, Lee CM, Rieger A, Fridgen G. Privacy-preserving federated learning for residential short-term load forecasting. *Appl Energy*. 2022;326:119915.
- Lin J, Ma J, Zhu J. Privacy-preserving household characteristic identification with federated learning method. *IEEE Trans Smart Grid*. 2021;13(2):1088-99.
- Fekri MN, Grolinger K, Mir S. Asynchronous adaptive federated learning for distributed load forecasting with smart meter data. *Int J Electr Power Energy Syst*. 2023;153:109285.
- Elgalhud MA, Fekri MN, Mir S, Grolinger K. Federated Online Learning for adaptive load forecasting across decentralized nodes. *Int J Electr Power Energy Syst*. 2025;169:110779.
- Haben S, Giasemidis G, Ziel F, Arora S. Short term load forecasting and the effect of temperature at the low voltage level. *Int J Forecast*. 2019;35(4):1469-84.
- Aurangzeb K, Haider SI, Alhussein M. Individual household load forecasting using bi-directional LSTM network with time-based embedding. *Energy Rep*. 2024;11:3963-75.
- Wang Y, Gao N, Hug G. Personalized federated learning for individual consumer load forecasting. *CSEE J Power Energy Syst*. 2022;9(1):326-30.

Kim JY, Cho SB. Electric energy consumption prediction by deep learning with state explainable autoencoder. *Energies*. 2019;12(4):739.

Tian Y, Sehovac L, Grolinger K. Similarity-based chained transfer learning for energy forecasting with big data. *IEEE Access*. 2019;7:139895-908.

Xu C, Chen G, Li C. Federated learning for interpretable short-term residential load forecasting in edge computing network. *Neural Comput Appl*. 2023;35(11):8561-74.

Zhu S, Shi X, Zhao H, Chen Y, Zhang H, Song X, et al. Personalized federated learning for household electricity load prediction with imbalanced historical data. *Appl Energy*. 2025;384:125419.

Rafi SH, Deeba SR, Hossain E. A short-term load forecasting method using integrated CNN and LSTM network. *IEEE access*. 2021;9:32436-48.

Mao Z, Li H, Huang Z, Yang C, Li Y, Zhou Z. Communication-efficient federated learning for power load forecasting in electric iots. *IEEE Access*. 2023;11:47930-9.

Manzoor HU, Khan AR, Flynn D, Alam MM, Akram M, Imran MA, et al. Fedbranched: Leveraging federated learning for anomaly-aware load forecasting in energy networks. *Sensors*. 2023;23(7):3570.

Dong Y, Wang Y, Gama M, Mustafa MA, Deconinck G, Huang X. Privacy-preserving distributed learning for residential short-term load forecasting. *IEEE Internet Things J*. 2024;11(9):16817-28.

Wang M, Xin R, Xia M, Zuo Z, Ge Y, Zhang P, et al. A federated LSTM network for load forecasting using multi-source data with homomorphic encryption. *AIMS Energy*. 2025;13(2):265-89.

Sarantinopoulos E, Michalakopoulos V, Sarmas E, Marinakis V, Todorean L, Cioara T. Meta-heuristic federated learning aggregation methods for load forecasting. *Energy AI*. 2025;22:100594.
<https://doi.org/10.1016/j.egyai.2025.100594>.

Todorean L, Daian M, Cioara T, Anghel I, Michalakopoulos V, Sarantinopoulos E, et al. Heuristic based federated learning with adaptive hyperparameter tuning for households energy prediction. *Sci Rep*. 2025;15(1):12564.

Sarker MA, Shanmugam B, Azam S, Thennadil S. Enhancing smart grid load forecasting: An attention-based deep learning model integrated with federated learning and XAI for security and interpretability. *Intell Syst Appl*. 2024;23:200422.

Richter L, Lenk S, Bretschneider P. Advancing electric load forecasting: Leveraging federated learning for distributed, non-stationary, and discontinuous time series. *Smart Cities*. 2024;7(4):2065-93.

Michalakopoulos V, Spyropoulos L, Ilias L, Sarmas E, Marinakis V, Askounis D. A comparative analysis of federated learning strategies for short-term photovoltaic power forecasting. *Energy Rep*. 2025;13:4223-35.

Zhou X, Feng J, Wang J, Pan J. Privacy-preserving household load forecasting based on non-intrusive load monitoring: A federated deep learning approach. *Peer J Comput Sci*. 2022;8:e1049.