

REVIEW

Open access

A Critical Review of Federated and Privacy-Preserving Learning for Smart-Meter Analytics, Grid-Edge Control, and DER Coordination

Nora Schneider^{1*}, Tobias Frank¹

Abstract

Smart-meter data is sensitive, and its centralization raises privacy concerns. Federated learning and privacy-preserving techniques promise to enable analytics without compromising user confidentiality. The practical viability of these methods in real-world smart-grid applications remains uncertain. A critical synthesis of their limitations, privacy assumptions, and deployment gaps is still needed. This review critically evaluates the state-of-the-art in federated and privacy-preserving learning for smart-meter analytics, grid-edge control, and DER coordination. It focuses on both methodological promise and deployment realism. A critical review of 27 peer-reviewed articles from 2017 to 2025 is conducted. The assessment focuses on technical performance, privacy guarantees, scalability, security, and deployment readiness. The review identifies key strengths in theoretical privacy and decentralized learning. It also finds substantial limitations in communication efficiency, non-IID robustness, adversarial resilience, and energy-sector-specific feasibility. Federated and privacy-preserving learning holds great promise, but critical technical and regulatory barriers must be overcome. The review provides a roadmap for future interdisciplinary research.

Keywords Critical review, Smart-meter analytics, Federated learning, Privacy-preserving learning, Grid-edge control, DER coordination

*Correspondence:

Nora Schneider
nora.schneider@gmail.com

¹ Department of Electrical Energy Systems, Faculty of Engineering, RWTH Aachen University, Aachen, Germany

Introduction

Smart-meter data provides fine-grained visibility into household and grid-edge behavior, but the same granularity makes it deeply privacy-sensitive. Differential privacy studies using real smart-meter data show that consumption traces can reveal behavioral patterns even before advanced learning models are applied [1]. The literature shows that this privacy challenge is not peripheral to energy analytics but structurally embedded in load forecasting, theft detection, demand response, and consumer profiling [2]. A critical review is therefore needed because privacy-preserving learning is often presented as a solution before its assumptions are tested under realistic grid conditions.

Federated learning has emerged as a leading approach because it allows models to be trained across distributed meters, homes, or utilities without directly pooling raw data. Early smart-grid applications used federated architectures to identify consumer characteristics while keeping data local, demonstrating the conceptual fit between decentralized learning and decentralized electricity use [2]. Residential load forecasting studies further show that federated recurrent and personalized models can exploit cross-household learning while reducing the need for centralized smart-meter

repositories [3, 4]. However, this review finds that decentralization alone does not guarantee privacy, robustness, or deployability.

Privacy-preserving smart-grid learning has also expanded beyond federated learning to include differential privacy, secure aggregation, secure multi-party computation, and homomorphic encryption. Differentially private optimal power flow and renewable-energy smart-metering studies indicate that formal privacy guarantees can be integrated into power-system optimization, but often with accuracy or operational-efficiency costs [5, 6]. Secure federated detection of false data injection attacks shows that privacy and security objectives may overlap, yet they can also compete when encrypted or noisy updates reduce situational awareness [7]. The literature therefore requires critical assessment rather than simple classification of techniques.

The need for review is amplified by the growing role of grid-edge intelligence, flexible demand, electric vehicles, storage, and distributed energy resources. Federated reinforcement learning has been proposed for smart-home energy management with distributed resources, suggesting that privacy-preserving learning is moving from offline analytics toward control [8]. More recent work on edge intelligence for smart meters illustrates the ambition to embed learning directly into distributed metering infrastructure [9]. Yet few studies fully address the latency, cyber-security, regulatory, and accountability requirements of real distribution-system operation.

Review Scope

This review covers federated learning, differential privacy, secure aggregation, secure multi-party computation, and homomorphic encryption when they are applied to smart-meter analytics, grid-edge control, or DER coordination. It includes load forecasting, consumer profiling, theft and anomaly detection, demand response baseline estimation, EV charging coordination, battery and flexible-load aggregation, and peer-to-peer energy trading. Studies on federated smart-meter forecasting, privacy-preserving theft detection, and federated data sharing are included because they directly address decentralized learning over sensitive energy data [3, 10, 11]. The review excludes purely centralized forecasting studies, non-peer-reviewed work, and methods that do not address privacy, decentralization, or grid-edge deployment.

The article is a critical review rather than a systematic meta-analysis or experimental benchmark. It evaluates how the selected literature frames privacy guarantees, performance trade-offs, scalability, communication efficiency, and real-world feasibility, including works on asynchronous federated learning, federated split learning, and heterogeneous federated learning [9, 12, 13]. Regulatory and compliance issues are considered only where they affect deployment readiness, data access, or energy-sector privacy governance. The review does not introduce new empirical results, simulated benchmarks, or new mathematical protocols.

Federated Learning Background

Federated learning paradigm

Federated learning trains a shared model across distributed clients by exchanging model parameters or updates instead of raw data, which makes it attractive for smart-meter environments. In residential load forecasting, this paradigm allows households or utility zones to contribute to a global model while retaining local consumption traces [3]. The literature shows that this architecture can reduce direct data centralization risks, but it does not remove privacy leakage from gradients, model updates, or client participation patterns [4]. A key limitation is that many energy studies still equate local data retention with complete privacy protection.

Privacy-preserving enhancements

Privacy-preserving enhancements attempt to address the leakage risks that remain after federated learning decentralizes data storage. Differential privacy introduces calibrated noise to limit the influence of individual records, and smart-grid

studies show its usefulness for both metering data and optimization settings [1, 5]. Secure aggregation and related cryptographic methods can hide individual updates from the server, while privacy-preserving federated theft-detection studies show that such protections are increasingly viewed as necessary rather than optional [14]. However, this review finds that stronger privacy guarantees often increase computational burden, reduce accuracy, or complicate deployment on constrained edge devices.

Relevance to distributed energy systems

Distributed energy systems are structurally aligned with federated learning because data, control authority, and physical assets are already dispersed across homes, feeders, aggregators, and devices. Federated reinforcement learning for multiple smart homes demonstrates how local energy agents can learn coordinated behavior without full data centralization [8]. Demand-response baseline estimation further illustrates the relevance of horizontal federated learning when customers or aggregators possess fragmented but complementary consumption records [15]. The critical challenge is that power systems require reliable and timely operation, while federated learning often assumes communication rounds and convergence conditions that may not match grid-control time scales.

Figure 1 illustrates how federated and privacy-preserving learning methods progress from distributed smart-meter and DER data environments toward operational deployment, while exposing the privacy, communication, robustness, scalability, and governance constraints identified in the review.




Figure 1. Privacy-Preserving Learning-to-Deployment Readiness Architecture for Smart-Meter Analytics, Grid-Edge Control, and DER Coordination

Privacy-Preserving Smart-Meter Analytics

Load forecasting and profiling

Load forecasting is one of the most mature applications of federated smart-meter analytics, with studies showing that recurrent, personalized, and asynchronous models can learn from distributed household data [3, 12]. Privacy-preserving

residential forecasting demonstrates clear value because individual consumption profiles can remain local while aggregated model knowledge improves prediction [4]. Personalized federated load prediction also addresses the fact that household demand is heterogeneous, but it introduces questions about fairness, model drift, and how much personalization weakens the benefits of federation [16]. The literature shows promise, yet most evaluations remain limited to historical datasets rather than live operational forecasting pipelines.

Anomaly and theft detection

Anomaly, false-data, and energy-theft detection studies show that federated learning is increasingly used for security-sensitive smart-grid analytics. FedDetect applies a privacy-preserving federated framework to energy-theft detection, indicating that distributed detection can reduce raw-data exposure while still identifying suspicious consumption behavior [10]. Secure federated deep learning for false data injection attacks extends this logic to cyber-physical grid security, where the detection target is not only consumer fraud but also malicious manipulation of grid measurements [7]. A key limitation is that adversaries may attack the federated process itself, which means privacy-preserving detection systems must also be robust against poisoning, evasion, and inference threats.

Consumer segmentation

Consumer segmentation and profiling highlight a tension at the center of privacy-preserving smart-meter analytics: the more useful the analytics, the more revealing the inferred categories may become. Federated consumer-characteristics identification shows that collaborative learning can classify or characterize electricity users without centralizing raw readings [2]. This is methodologically important, but it also raises a critical ethical issue because privacy risk can shift from raw data exposure to inferred behavioral labels. The review therefore finds that privacy-preserving profiling must be assessed not only by data-locality claims but also by the sensitivity of the outputs it produces.

Comparative summary of techniques

The comparative literature indicates that no single privacy-preserving technique dominates across accuracy, overhead, and privacy guarantee. Differential privacy provides formal guarantees but can degrade performance when noise is high, especially in fine-grained metering or control applications [1]. Federated learning reduces centralized data collection but remains vulnerable unless combined with secure aggregation, robust aggregation, or privacy accounting, as shown by privacy-preserving theft-detection and communication-efficient prediction schemes [14, 17]. **Table 1** summarizes the main privacy techniques used or implied in the reviewed smart-meter analytics literature and highlights their practical trade-offs.

Table 1. Privacy Techniques for Smart-Meter Analytics

Privacy technique	Typical smart-meter use	Accuracy impact	Computational overhead	Privacy guarantee	Critical assessment
Federated learning	Load forecasting, profiling, theft detection, false-data detection	Usually moderate when data are sufficiently similar, but weaker under non-IID household behavior	Moderate communication and client-training burden	Keeps raw data local but does not by itself prevent update leakage	Strong practical appeal, but privacy is often overstated when no secure aggregation or inference defense is used

Differential privacy	Smart-meter data release, demand response, optimal power flow, privacy-aware analytics	Can be significant when strict privacy budgets require high noise	Low to moderate computational cost, but tuning privacy budgets is difficult	Formal record-level or user-level privacy depending on mechanism design	Strong theoretical protection, but utility loss and unclear energy-sector privacy-budget norms limit deployment
Secure multi-party computation	Collaborative analytics or aggregation without revealing individual inputs	Usually low direct accuracy loss because data are protected cryptographically rather than noised	High communication and synchronization overhead	Strong cryptographic protection under specified threat models	Powerful for aggregation, but often difficult to scale to millions of meters and low-latency control
Homomorphic encryption	Encrypted inference or aggregation over sensitive metering inputs	Usually low direct accuracy loss, though approximations may affect complex models	High computational cost, especially for deep learning and edge devices	Allows computation on encrypted data under cryptographic assumptions	Attractive for high-sensitivity use cases, but currently constrained by latency, device capability, and operational complexity

Grid-Edge Control Applications

Voltage and frequency control

Grid-edge control requires analytics to move from retrospective prediction toward decisions that affect voltage, frequency, and local balancing. Differentially private optimal power flow shows that privacy-preserving methods can be embedded in grid optimization, but it also demonstrates that privacy protection may alter solution quality and operational feasibility [5]. Risk assessment based on federated learning extends the role of privacy-preserving AI toward operational decision support, where distributed data can inform grid risk without full centralization [18]. Few studies, however, validate whether these methods can satisfy the reliability expectations of real-time distribution automation.

Real-time edge inference

Real-time edge inference is essential because smart meters, gateways, and local controllers cannot always wait for centralized cloud processing. Federated split learning for smart meters directly targets this problem by distributing intelligence between edge devices and higher-level infrastructure [9]. The literature shows that such architectures may reduce bandwidth and protect raw data, but they also introduce dependencies between device capacity, partitioned model design, and communication stability. A key limitation is that many studies demonstrate inference feasibility without fully evaluating maintenance, software updates, fault recovery, and cyber-physical safety.

Communication constraints

Communication is one of the strongest practical constraints on federated learning at the grid edge. Communication-efficient energy prediction schemes recognize that repeated model-update exchange can burden smart-grid communication networks, especially when the number of meters or clients grows [17]. Asynchronous adaptive federated learning attempts to reduce synchronization bottlenecks in distributed load forecasting, which is important because meters and edge devices may be intermittently connected or resource constrained [12]. This review finds that

communication efficiency is often treated as a model-training issue, while its interaction with utility network operations and service-level requirements remains underexplored.

Privacy-utility trade-offs

Grid-edge control exposes privacy-utility trade-offs more sharply than offline analytics because control decisions must be accurate, timely, and explainable. Privacy-preserving false-data detection can protect distributed measurements, but excessive noise, encryption latency, or incomplete client participation may weaken situational awareness [19]. Federated energy management in smart homes shows that decentralized learning can support privacy-aware control, yet performance depends on the representativeness of local agents and the stability of the environment [8]. The literature therefore supports privacy-preserving grid-edge intelligence in principle, but real-world feasibility depends on explicit accounting for latency, cyber-security, regulatory accountability, and operational fallback mechanisms.

DER Coordination Applications

EV charging coordination

EV charging coordination is a natural application for federated and privacy-preserving learning because charging profiles reveal mobility, occupancy, and socioeconomic patterns. Federated prediction of EV charging load shows that decentralized learning can support charging-demand estimation while reducing the need to centralize user-level charging records [20]. However, this review finds that prediction alone is insufficient for coordination because operational charging control also requires constraints, feeder limits, tariffs, and user preferences to be represented reliably. A key limitation is that few EV-focused studies connect federated learning outputs to validated real-time charging decisions under distribution-network constraints.

Battery storage aggregation

Battery storage aggregation requires coordination across many small devices whose state-of-charge, availability, and cycling preferences may be commercially and personally sensitive. Federated energy optimization frameworks suggest that local devices can participate in shared learning objectives without directly exposing operational data, which is relevant for aggregating behind-the-meter storage [21]. Personalized federated household load prediction further indicates that device-level heterogeneity cannot be ignored when batteries are coordinated through learned forecasts [22]. The literature shows promise, but it rarely evaluates battery degradation, aggregator incentives, or conflicting objectives between customer value and grid services.

Demand response with privacy

Demand response is one of the strongest cases for privacy-preserving DER coordination because it depends on flexible consumption while exposing detailed household behavior. Adaptive horizontal federated learning for demand-response baseline estimation demonstrates that privacy-aware learning can support customer-specific baselines without fully centralizing consumption histories [15]. Differential privacy for renewable-energy smart metering also shows that privacy mechanisms can be connected to distributed energy participation, although utility loss remains a central concern [6]. This review finds that demand-response studies often understate the difficulty of verifying performance, detecting gaming, and allocating incentives when the underlying data remain distributed or perturbed.

Multi-agent coordination

Multi-agent DER coordination extends privacy-preserving learning from prediction to strategic interaction among homes, aggregators, prosumers, and market platforms. Federated reinforcement learning for decentralized peer-to-peer energy trading shows how agents can learn trading behavior without centralizing all participant data [23]. Integrated blockchain and federated learning for DER cybersecurity further suggests that privacy-preserving coordination may need tamper-

resistance, auditability, and cyber-defense mechanisms beyond model training alone [24]. A key limitation is that multi-agent learning can amplify instability, strategic manipulation, and fairness concerns when participants have unequal data quality, device flexibility, or market power.

Critical Limitations

Communication overhead and latency

Communication overhead remains one of the most persistent limitations in federated smart-grid learning. Residential federated load forecasting can require repeated exchange of model updates, and even efficient implementations may become burdensome as client numbers increase [25]. Communication-efficient energy prediction explicitly addresses this issue, but its optimization of message exchange does not fully solve latency requirements for protection, control, or fast DER dispatch [17]. This review finds that many studies evaluate communication cost as bandwidth or rounds while giving less attention to utility communication reliability, outage conditions, and operational prioritization.

Non-IID data and model performance

Non-IID data are not an exception in smart-meter analytics but the normal condition because households differ in appliances, occupancy, climate exposure, tariffs, and DER ownership. Personalized federated learning for consumer load forecasting shows that customization can improve performance when generic global models fail to capture individual patterns [16]. Heterogeneous federated learning for electricity theft detection similarly shows that class imbalance and client heterogeneity must be handled explicitly rather than treated as minor dataset imperfections [13]. A key limitation is that personalization can reduce comparability, complicate privacy accounting, and make system-wide validation harder.

Security vulnerabilities (backdoors, inference)

Federated learning can reduce raw-data exposure while creating new attack surfaces through model updates, aggregation protocols, and client participation. A privacy-preserving federated learning scheme against poisoning attacks shows that poisoning is a serious concern in smart-grid deployments where compromised clients may manipulate model behavior [26]. Secure false-data injection detection studies highlight that cyber-security and privacy must be jointly considered because attackers may target both the grid measurements and the learning pipeline [7]. This review finds that few studies provide end-to-end threat models covering inference attacks, poisoning, backdoors, compromised aggregators, and malicious edge devices simultaneously.

Lack of standardized benchmarks

A major weakness of the literature is the lack of standardized benchmarks that jointly evaluate forecasting accuracy, privacy leakage, communication cost, robustness, and deployment feasibility. Studies on smart-meter data sharing, theft detection, forecasting, and risk assessment often use different datasets, privacy assumptions, and evaluation metrics, which makes cross-study comparison difficult [11]. Broader privacy-preservation reviews for smart meters emphasize that the field lacks consistent ways to compare privacy mechanisms under realistic household and grid conditions [27]. **Table 2** synthesizes the critical limitations identified across the reviewed literature and links them to representative references.

Table 2. Critical Limitations and Associated Literature

Critical limitation	How it appears in the literature	Why it matters for deployment	Representative references
Communication cost	Repeated federated rounds, client synchronization, and update	Large-scale smart-meter fleets may not support frequent model exchange	[12, 17, 25]

	transmission burden smart-grid communication networks	during congestion, outages, or low-bandwidth operation	
Non-IID and imbalanced data	Household behavior, theft labels, EV charging, and DER flexibility differ strongly across clients	Global models may underperform for minority users, unusual feeders, or rare events, weakening reliability and fairness	[13, 16, 20, 22]
Privacy leakage through updates	Local data retention does not prevent leakage through gradients, participation patterns, or model outputs	Utilities may still face privacy risk even when raw smart-meter data are not centralized	[1, 4, 14, 19]
Adversarial robustness	Poisoning, false-data injection, compromised clients, and cyber-physical manipulation can target federated pipelines	Privacy-preserving learning may fail if attackers corrupt model updates or exploit trusted aggregation assumptions	[7, 24, 26]
Computation at the edge	Smart meters and gateways may have limited processing, memory, and energy budgets	Cryptographic protection, split learning, and deep models may be difficult to maintain at scale	[5, 6, 9]
Benchmark inconsistency	Studies use different datasets, threat models, metrics, and privacy definitions	Claims of superiority are difficult to compare, reproduce, or translate into procurement and regulation	[11, 18, 27]

Privacy/Deployment Gaps

Regulatory and compliance gaps

Regulatory compliance remains underdeveloped in much of the technical literature, even when privacy is the stated motivation. Differential privacy for smart-metering provides formal protection concepts, but translating privacy budgets into regulatory acceptability is still unresolved [1]. Privacy preservation reviews highlight that smart-meter data governance involves consent, access control, retention, and secondary use, not only model-training architecture [27]. This review finds that federated learning studies often do not specify who controls aggregation, who audits privacy claims, or how consumers can exercise data rights.

Operational barriers in real grids

Operational deployment faces barriers that are not captured by offline simulations or historical-data experiments. Federated split learning for smart meters points toward practical edge intelligence, yet maintaining models across heterogeneous meter generations, firmware environments, and communication links remains a difficult operational task [9]. Federated risk assessment shows the value of distributed learning for grid decision support, but real operators also need explainability, fallback procedures, and integration with supervisory control systems [18]. Few studies address maintenance responsibility, incident response, or liability when a privacy-preserving model contributes to an operational decision.

Scalability to millions of meters

Scalability to millions of meters is not merely a larger version of laboratory-scale federation. Communication-efficient prediction schemes and asynchronous federated learning reduce some scaling pressure, but both still require careful orchestration of participation, update timing, and model-version control [12, 17]. Federated smart-meter data sharing

raises additional scalability concerns because synthetic or shared representations may still require privacy validation across diverse customer populations [11]. **Table 3** maps the main privacy and deployment gaps to regulatory, technical, and operational barriers, while proposing directions that are consistent with the evidence reviewed.

Table 3. Privacy and Deployment Gaps with Proposed Solutions

Gap	Deployment barrier	Barrier type	Proposed solution
Undefined privacy budgets for energy data	Utilities and regulators lack shared thresholds for acceptable privacy-utility trade-offs	Regulatory	Develop sector-specific privacy-budget guidance for metering, forecasting, demand response, and DER coordination
Weak governance of aggregation roles	Aggregators, utilities, vendors, and platform operators may have unclear responsibilities	Regulatory	Require auditable role definitions, data-processing agreements, and independent privacy audits
Limited edge-device capability	Smart meters may not support heavy cryptography, deep models, or frequent updates	Technical	Use lightweight models, split learning, tiered aggregation, and hardware-aware privacy mechanisms
Incomplete adversarial threat models	Studies often isolate poisoning, inference, or false-data attacks rather than combining them	Technical	Build unified threat models covering clients, aggregators, communication links, and cyber-physical impacts
Poor integration with utility operations	Research prototypes rarely connect to control rooms, outage management, or DER management systems	Operational	Conduct field pilots with operator-in-the-loop evaluation, fallback procedures, and incident-response protocols
Lack of comparable benchmarks	Claims are difficult to reproduce across datasets and privacy assumptions	Operational	Establish shared benchmark suites reporting accuracy, privacy leakage, communication cost, latency, robustness, and fairness

Future Research

Communication-efficient FL architectures

Future research should prioritize communication-efficient federated architectures designed specifically for utility communication networks rather than generic distributed machine-learning environments. Asynchronous adaptive federated learning is a useful step because it reduces strict synchronization assumptions in distributed load forecasting [12]. Communication-efficient energy prediction also shows that model exchange can be optimized, but future designs need event-triggered updates, hierarchical aggregation, and feeder-aware participation policies [17]. This review finds that communication efficiency should be evaluated alongside operational latency, outage resilience, and cyber-security rather than as a standalone machine-learning metric.

Hybrid privacy models

Hybrid privacy models are likely to be more realistic than reliance on any single privacy-preserving technique. Federated learning can reduce raw-data centralization, differential privacy can bound individual contribution leakage, and cryptographic aggregation can protect updates during collaboration [14]. Differentially private optimal power flow and renewable-energy smart-metering studies show that formal privacy can be integrated into energy optimization, but utility

loss must be explicitly managed [5, 6]. Future work should therefore develop layered privacy architectures with transparent privacy accounting, attack testing, and adaptive protection levels for analytics, control, and DER coordination.

Regulatory sandboxes and field trials

The most important next step is movement from retrospective datasets to regulated field trials and sandboxes. Federated reinforcement learning for smart-home energy management and decentralized peer-to-peer energy trading show that privacy-preserving learning can support active coordination, but these approaches need validation under real tariffs, customer behavior, and operational constraints [8, 23]. Cyber-secure DER coordination research further indicates that privacy, auditability, and resilience must be tested together before deployment [24]. This review finds that interdisciplinary pilots involving utilities, regulators, vendors, aggregators, and consumers are essential to determine whether privacy-preserving learning can be trusted in real energy systems.

Conclusion

This critical review finds that federated and privacy-preserving learning has become a credible research direction for smart-meter analytics, grid-edge intelligence, and DER coordination. The strongest evidence appears in load forecasting, theft detection, privacy-aware prediction, and emerging edge-intelligence architectures. These studies show that decentralizing data access can reduce direct exposure of sensitive consumption records. However, the field is still far from proving full operational readiness.

The central weakness is that privacy is often framed too narrowly. Keeping raw data local is important, but it does not automatically prevent leakage through model updates, inferred labels, participation behavior, or compromised clients. Stronger privacy mechanisms can improve protection, yet they frequently introduce accuracy loss, computational cost, latency, and governance complexity. The practical question is therefore not whether privacy-preserving learning works in principle, but whether it works safely, reliably, and accountably in real grid environments.

The review also finds that DER coordination raises harder challenges than offline smart-meter analytics. EV charging, flexible demand, batteries, and peer-to-peer trading require timely decisions, fair incentives, cyber-resilience, and compatibility with distribution-network limits. These requirements expose gaps in benchmarking, threat modeling, regulatory interpretation, and operational integration. Future progress depends on evaluating privacy-preserving learning as part of the energy system, not only as a machine-learning protocol.

Federated and privacy-preserving learning should therefore be advanced through collaborative field pilots, standardized benchmarks, transparent privacy accounting, and sector-specific regulation. Utilities, researchers, technology providers, regulators, and consumer representatives need shared methods for testing privacy, accuracy, robustness, fairness, and control performance. Without this coordination, the field risks producing elegant algorithms that remain difficult to deploy. With it, privacy-preserving intelligence can become a foundation for trustworthy digital energy systems.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 26 Sep 2024

Revised: 03 Nov 2024

Accepted: 12 Dec 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

- Eibl G, Engel D. Differential privacy for real smart metering data. *Comput Sci Res Dev*. 2017;32(1):173-82.
- Wang Y, Bennani IL, Liu X, Sun M, Zhou Y. Electricity consumer characteristics identification: A federated learning approach. *IEEE Trans Smart Grid*. 2021;12(4):3637-47.
- 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.
- 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.
- Yang Z, Cheng P, Chen J. Differential-privacy preserving optimal power flow in smart grid. *IET Gener Transm Distrib*. 2017;11(15):3853-61.
- 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.
- Li Y, Wei X, Li Y, Dong Z, Shahidehpour M. Detection of false data injection attacks in smart grid: A secure federated deep learning approach. *IEEE Trans Smart Grid*. 2022;13(6):4862-72.
- Lee S, Choi DH. Federated reinforcement learning for energy management of multiple smart homes with distributed energy resources. *IEEE Trans Ind Inform*. 2020;18(1):488-97.
- Li Y, Qin D, Poor HV, Wang Y. Introducing edge intelligence to smart meters via federated split learning. *Nat Commun*. 2024;15(1):9044.
- Wen M, Xie R, Lu K, Wang L, Zhang K. FedDetect: A novel privacy-preserving federated learning framework for energy theft detection in smart grid. *IEEE Internet Things J*. 2021;9(8):6069-80.
- Chen Z, Li J, Cheng L, Liu X. Federated-WDCGAN: A federated smart meter data sharing framework for privacy preservation. *Appl Energy*. 2023;334:120711.

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.

Wen H, Liu X, Lei B, Yang M, Cheng X, Chen Z. A privacy-preserving heterogeneous federated learning framework with class imbalance learning for electricity theft detection. *Appl Energy*. 2025;378:124789.

Ashraf MM, Waqas M, Abbas G, Baker T, Abbas ZH, Alasmay H. Feddp: A privacy-protecting theft detection scheme in smart grids using federated learning. *Energies*. 2022;15(17):6241.

Wang R, Qiu H, Gao H, Li C, Dong ZY, Liu J. Adaptive horizontal federated learning-based demand response baseline load estimation. *IEEE Trans Smart Grid*. 2023;15(2):1659-69.

Wang Y, Gao N, Hug G. Personalized federated learning for individual consumer load forecasting. *CSEE J Power Energy Syst*. 2022;9(1):326-30.

Badr MM, Mahmoud MM, Fang Y, Abdulaal M, Aljohani AJ, Alasmay W, et al. Privacy-preserving and communication-efficient energy prediction scheme based on federated learning for smart grids. *IEEE Internet Things J*. 2023;10(9):7719-36.

Deng S, Zhang L, Yue D. Data-driven and privacy-preserving risk assessment method based on federated learning for smart grids. *Commun Eng*. 2024;3(1):154.

Tran HY, Hu J, Yin X, Pota HR. An efficient privacy-enhancing cross-silo federated learning and applications for false data injection attack detection in smart grids. *IEEE Trans Inf Forensics Secur*. 2023;18:2538-52.

Yin W, Ji J. Prediction of EV charging load based on federated learning. *Energy*. 2025;316:134559.

Gupta H, Agarwal P, Gupta K, Baliarsingh S, Vyas OP, Puliafito A. Fedgrid: A secure framework with federated learning for energy optimization in the smart grid. *Energies*. 2023;16(24):8097.

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.

Ye Z, Qiu D, Li S, Fan Z, Strbac G. Federated reinforcement learning for decentralized peer-to-peer energy trading. *Energy AI*. 2025;20:100500.

Babayomi OO, Igboanusi IS, Ahakonye LA, Kim DS. Integrated blockchain and federated learning for the cybersecurity of distributed energy resources. *Int J Electr Power Energy Syst*. 2025;173:111286.

Briggs C, Fan Z, Andras P. Federated learning for short-term residential load forecasting. *IEEE Open Access J Power Energy*. 2022;9:573-83.

Li X, Wen M, He S, Lu R, Wang L. A privacy-preserving federated learning scheme against poisoning attacks in smart grid. *IEEE Int Things J*. 2024;11(9):16805-16.

Kua J, Hossain MB, Natgunanathan I, Xiang Y. Privacy preservation in smart meters: current status, challenges and future directions. *Sensors*. 2023;23(7):3697.