Engineering and Technology Journal e-ISSN: 2456-3358

Volume 11 Issue 09 September-2026, Page No.- 11549-11560

DOI: 10.47191/etj/v11i09.19, I.F. – 8.595

Reducing Residential Grid Congestion from Electric Vehicles Using Privacy-Preserving Federated Smart Charging

Omar Ahmed Mohamed Edbeib[1*], Abdussalam Ali Ahmed[2]

[1,2]Mechanical and Industrial Engineering Department, Bani Waleed University, Libya

ABSTRACT: Residential electric vehicle charging can create short feeder peaks during already crowded evening hours. Utilities need coordinated charging, yet household charging records reveal daily routines and travel habits. This paper develops a privacy-preserving federated smart charging method for residential feeders. Each client trains an energy predictor without sending raw sessions to a server. Clipped updates receive distributed Gaussian noise before pairwise-mask secure aggregation. A feeder-aware scheduler then places predicted energy within each vehicle connection window. The experiment uses 6,163 measured charging sessions from 74 apartment users. It also uses measured household imports from Open Power System Data. The corrected workflow converts cumulative meter readings into hourly demand increments. Testing covers 958 unseen sessions on a modeled feeder serving 100 homes. The proposed private model uses a noise multiplier of 1.0. Its privacy budget is 14.91 at a delta value of 0.00001. Private smart charging lowers the combined peak from 160.07 kW to 100.55 kW. This reduction equals 37.18 percent and removes ten modeled overload hours. However, the schedule leaves 10.57 percent of requested energy unmet. Its mean charging delay reaches 8.37 hours. The private predictor records 12.80 kWh test RMSE. The centralized model performs better, with 8.60 kWh RMSE. Results show that privacy noise adds little error beyond ordinary federated training. Client data differences remain the main accuracy barrier. The method protects data location and bounds shared updates. It exposes the service cost behind congestion relief. These findings support stronger local models and service constraints before deployment.

KEYWORDS: electric vehicles, federated learning, smart charging, differential privacy, secure aggregation, residential feeder.

1. INTRODUCTION

Electric vehicles move transport demand from fuel stations into local electricity networks. Most private cars remain connected for longer than their batteries need. That unused time creates valuable charging flexibility. It also creates a control problem for residential feeders. Many drivers arrive after work and start charging immediately. Their demand can overlap with cooking, heating, cooling, and other household loads. Uncoordinated charging can therefore create sharp local peaks. These peaks may exceed transformer limits despite moderate daily energy use.

Earlier grid studies show that charging time matters as much as electric vehicle adoption. Muratori (2018) found that uncoordinated home charging can reshape residential demand. Clement-Nyns et al. (2010) reported higher losses and voltage stress under clustered charging. Transformer aging can also accelerate when charging repeats during local peaks (Hilshey et al., 2013). These risks do not require every household to own a vehicle. They can appear first on weak feeders with similar travel schedules. Targeted charging control is therefore more practical than broad network reinforcement.

Smart charging can move energy toward quieter hours while respecting vehicle departure times. Classical methods often require detailed sessions at a central controller. Those records contain plug-in times, energy demand, and repeated user identifiers. They can reveal work patterns, holidays, and periods away from home. Central collection therefore creates a privacy and governance burden. It also creates one attractive location for misuse or data theft. A useful charging method should limit raw data exposure from its first design stage.

Federated learning offers one route toward that goal. Clients train a shared model while their examples remain locally stored (McMahan et al., 2017). However, federated learning alone does not ensure formal privacy. Model updates can still carry information about local records (Kairouz et al., 2021). Secure aggregation hides each update inside a combined sum (Bonawitz et al., 2017). Differential privacy further limits the effect of any participating client. Combining these mechanisms can support useful learning with measurable disclosure limits.

This paper tests that combined design with measured charging and household data. The method predicts session energy through federated client training. It clips each local update before adding distributed Gaussian noise. Pairwise masks then emulate secure aggregation at the server boundary. A scheduler uses the private prediction and a calibration margin. It assigns charging to lower-load hours inside each connection window. The evaluation reports prediction, congestion, privacy, and service results together. This joint reporting prevents peak reduction from hiding unmet driver needs.

The paper makes three connected contributions. First, it presents a reproducible client-level private training workflow for charging prediction. Second, it links that predictor to a capacity-aware residential scheduler. Third, it measures the resulting tension between feeder relief and charging service. The evidence comes from a corrected experiment with chronological client splits. Cumulative household meters are converted into hourly increments before feeder construction. The findings show strong congestion relief, but they also identify an important service penalty.

2. LITERATURE REVIEW

2.1 Residential charging and grid congestion

Residential charging creates a local timing problem rather than only an energy problem. A vehicle may consume modest daily energy but draw several kilowatts continuously. Several simultaneous chargers can dominate a small feeder during narrow evening windows. Richardson (2013) linked these patterns with generation, network, and market impacts. Muratori (2018) later used detailed residential demand to show nonlinear peak growth. That result matters because average demand can hide short overload periods. Feeder studies must therefore report peaks, duration, and background demand together.

Distribution equipment responds to both load size and load duration. Hilshey et al. (2013) connected smart charging with distribution transformer aging. Their work showed how timing control can reduce thermal stress. Clement-Nyns et al. (2010) also found power losses and voltage effects. Those studies support capacity-aware charging at the feeder level. They also warn against evaluating charging only through energy prices. A low-cost schedule may still concentrate demand on a constrained neighborhood asset.

2.2 Coordinated and decentralized smart charging

Coordinated charging has been studied through optimization, pricing, and decentralized control. Sortomme and El-Sharkawi (2011) formulated optimal unidirectional charging strategies. Gan et al. (2013) developed a decentralized protocol with feeder objectives. Both approaches show that charging flexibility can reduce harmful peaks. Recent reviews describe many objectives, including cost, voltage, and renewable use (Sadeghian et al., 2022). Yet practical methods need forecasts for uncertain session energy and departure behavior. Forecast errors can cause missed energy even when the optimization appears feasible.

Centralized control offers direct coordination, but it concentrates private records. Decentralized optimization reduces that concentration, though messages may still reveal behavior. The privacy problem therefore extends beyond raw database storage. Repeated gradients, bids, or flexibility reports can expose patterns. A private design must protect the messages leaving each client. It must also quantify the privacy loss across repeated communication rounds. These needs lead naturally toward federated learning with formal privacy accounting.

2.3 Federated learning and client heterogeneity

Federated averaging combines locally trained model changes through weighted aggregation. McMahan et al. (2017) showed its communication benefits across decentralized data. Charging data, however, are strongly non-identical across households. Some clients charge daily, while others charge rarely or irregularly. Energy needs also differ by vehicle, commute, season, and charger access. These differences can pull local models toward conflicting targets. Kairouz et al. (2021) identify such statistical heterogeneity as a central federated challenge.

FedAvg can lose accuracy when local data distributions differ strongly. Li et al. (2020) proposed FedProx to improve training under heterogeneous conditions. Their analysis explains why ordinary averaging may drift during repeated local updates. Charging applications face the same issue because each client holds a personal history. A private deployment also restricts central inspection of that history. The controller cannot easily repair poorly represented clients through direct pooling. Prediction comparisons must therefore include centralized and nonprivate federated baselines.

2.4 Differential privacy and secure aggregation

Differential privacy limits how much one participant can change an observed result. Dwork and Roth (2014) provide the main mathematical foundation. Training usually clips individual contributions before adding calibrated noise. Abadi et al. (2016) demonstrated practical accounting for noisy gradient training. Rényi differential privacy supports tighter composition across many steps (Mironov, 2017). The final privacy budget depends on clipping, noise, sampling, and communication rounds. A reported private model should therefore include epsilon and delta values.

Secure aggregation solves a different problem. It prevents the server from reading a single client update during aggregation. Bonawitz et al. (2017) designed a practical protocol for this setting. Differential privacy remains useful because the final aggregate can still reveal information. The two protections therefore complement each other. This paper combines clipped noisy updates with a pairwise-mask aggregation emulation. The emulation validates arithmetic, but it is not a production cryptographic deployment.

2.5 Research gap

Much existing work studies charging control, federated prediction, or privacy separately. That separation can hide practical trade-offs. An accurate private model may still create poor driver service. A smooth feeder profile may also depend on unrealistic background load data. This paper joins measured sessions, measured household imports, private learning, and scheduling. It compares model error with grid relief and unmet energy. The design therefore evaluates the complete path from local records to feeder operation.

Table 1: Position of the present work within related research

Study	Main focus	Method	Remaining issue
Clement-Nyns et al. (2010)	Grid impact	Load flow study	No formal privacy
Gan et al. (2013)	Peak control	Decentralized optimization	No learned demand model
Muratori (2018)	Residential peaks	Measured demand simulation	No privacy mechanism
Li et al. (2020)	Federated heterogeneity	FedProx	Not a charging study
Sadeghian et al. (2022)	Smart charging review	Literature synthesis	No joint experiment
Present paper	Private feeder control	DP-FedAvg and scheduling	Reports service trade-off

3. PROPOSED METHOD

3.1 System structure

The modeled system contains residential clients, an aggregation server, and a feeder scheduler. Each client stores its own charging sessions. A session includes arrival, departure, energy, user type, and calendar variables. The client trains a common neural network from its local training records. Only protected parameter updates leave the client. The server produces one global energy predictor. That predictor supports charging allocation under a shared feeder capacity.

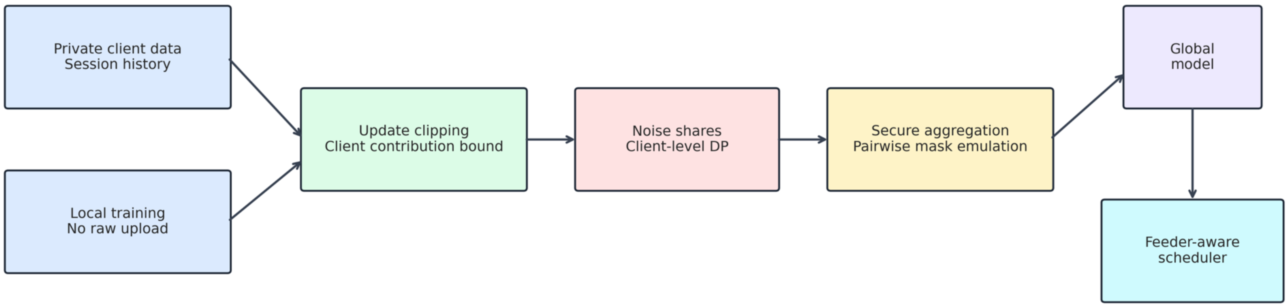

Figure 1 Privacy-preserving federated smart charging framework.

Figure 1 shows the information boundary used in the experiment. The server never receives raw charging examples. Clients clip their updates to a fixed Euclidean norm. They then add independent Gaussian noise shares. Pairwise masks cancel after the updates are summed. The server observes only the protected aggregate. This sequence separates data location, message secrecy, and formal privacy.

3.2 Local energy prediction

The prediction target is charged energy for one session. Fourteen features describe time, connection length, user type, and local history. Cyclic sine and cosine terms represent hour, weekday, and month. A logarithmic feature represents connection duration. Two prequential features describe earlier energy use by the same client. Those features use only sessions that occurred before the current record. This rule prevents future information from entering model training.

The predictor is a small multilayer perceptron. It contains hidden layers with 32 and 16 rectified linear units. A final linear unit estimates normalized log energy. The model uses Huber loss because rare large sessions can distort squared error. Predictions return to kilowatt-hours through the inverse log transformation. They are clipped between zero and 80 kWh. The same architecture supports centralized, federated, and private comparisons.

$$\hat{y} = \exp(\sigma y\, f\theta(x) + \mu y) - 1$$

Here, fθ maps the feature vector into normalized log energy. The symbols μy and σy are training-set moments. Only these two aggregate moments are shared across clients. A small network keeps local training practical for residential gateways. The experiment does not measure gateway runtime or communication energy. Those costs require field hardware and network traces.

3.3 Federated averaging

The global model begins from a seeded common parameter vector. Each round samples 30 percent of the 74 eligible clients. This rule selects 23 clients per round. Every selected client trains for two local epochs. The local optimizer uses Adam with a learning rate of 0.025. The client returns its parameter change rather than its examples. The server averages selected changes and updates the global model.

$$\theta t + 1 = \theta t + \left(\frac{1}{m}\right) \Sigma k \in St\ \Delta\theta k, t$$

The symbol St denotes selected clients during round t. The value m is the number of selected clients. Ordinary FedAvg uses 40 rounds in the full experiment. Centralized training uses 100 epochs on pooled training records. Both approaches share the same features and neural architecture. This controls architecture differences during the accuracy comparison.

3.4 Client-level differential privacy

Each client update is clipped before aggregation. The clipping norm equals 0.50 across every private experiment. Clipping limits the maximum influence of one participating client. Independent clients then add Gaussian noise shares. Those shares sum to the intended aggregate mechanism before server observation. The main configuration uses a noise multiplier of 1.0. Additional runs use multipliers of 0.5 and 1.5.

$$\Delta\theta k = \Delta\theta k \cdot min(1, C / ||\Delta\theta k||2)$$

$$\theta t + 1 = \theta t + \left(\frac{1}{m}\right)[\Sigma k \in St\, \Delta\bar{\theta} k + N(0, \sigma^2 C^2 I)]$$

The clipping bound is C, while sigma controls noise strength. Privacy is tracked through a Rényi accountant (Mironov, 2017). The reported budget uses a delta value of 0.00001. Larger noise generally lowers epsilon and strengthens the stated guarantee. However, noise can reduce prediction utility. The experiment measures that trade-off on the unchanged test split.

3.5 Secure aggregation emulation

The experiment emulates pairwise masking between every selected client pair. One client adds a random mask, while its paired client subtracts that mask. All masks cancel when the server sums the messages. The resulting total should match the direct protected sum. The maximum observed difference provides an arithmetic check. This procedure follows the secure aggregation principle from Bonawitz et al. (2017). It omits encrypted transport, authentication, dropout recovery, and malicious-client defenses.

3.6 Feeder-aware charging scheduler

The scheduler operates on hourly slots across every test connection window. A 7.2 kW charger sets the maximum hourly charging rate. The scheduler orders sessions by departure time and remaining slack. It then assigns requested energy to the lowest reserved feeder-load slots. Assignments never exceed the modeled feeder capacity. A rolling check releases unused power when the actual battery stops early. The private request equals predicted energy plus a validation margin.

$$Lt = Bt + \Sigma i\, pi,t \text{ and } Lt \le Cfeeder$$

The value Bt is measured background demand at hour t. The term pi,t is charging power for session i. The scheduler minimizes peaks through a lowest-load placement rule. It does not solve a full mixed-integer optimization problem. This choice keeps the experiment transparent and reproducible. It also allows prediction error to appear directly as unmet energy.

3.7 Evaluation measures

Prediction quality uses mean absolute error, root mean square error, and R squared. A validation error quantile creates the upper calibration margin. Grid quality uses combined peak, peak reduction, load variance, and peak-to-average ratio. Overload hours count hourly demand above feeder capacity. Service quality uses unmet energy and energy-weighted charging delay. Privacy reporting includes epsilon, delta, clipping behavior, and secure-sum error. Together, these measures show whether congestion relief shifts costs onto drivers.

4. DATA AND EXPERIMENTAL DESIGN

4.1 Electric vehicle charging sessions

Charging records come from apartment buildings in Norway. The public dataset includes charging reports and repeated user identifiers. Sørensen et al. (2021a) describe its collection and release. A related study examines residential charging flexibility from these records (Sørensen et al., 2021b). The raw charging report contains 6,878 rows. Cleaning removes missing values and implausible energy or duration values. Clients also need at least 15 chronological records.

The final sample contains 6,163 sessions from 74 users. Their combined charged energy equals 84,202.1 kWh. Median session energy is 9.9 kWh. Median connection duration is 11.0 hours. These values confirm that many sessions contain scheduling slack. However, session counts vary strongly across users. That imbalance creates a realistic federated learning challenge.

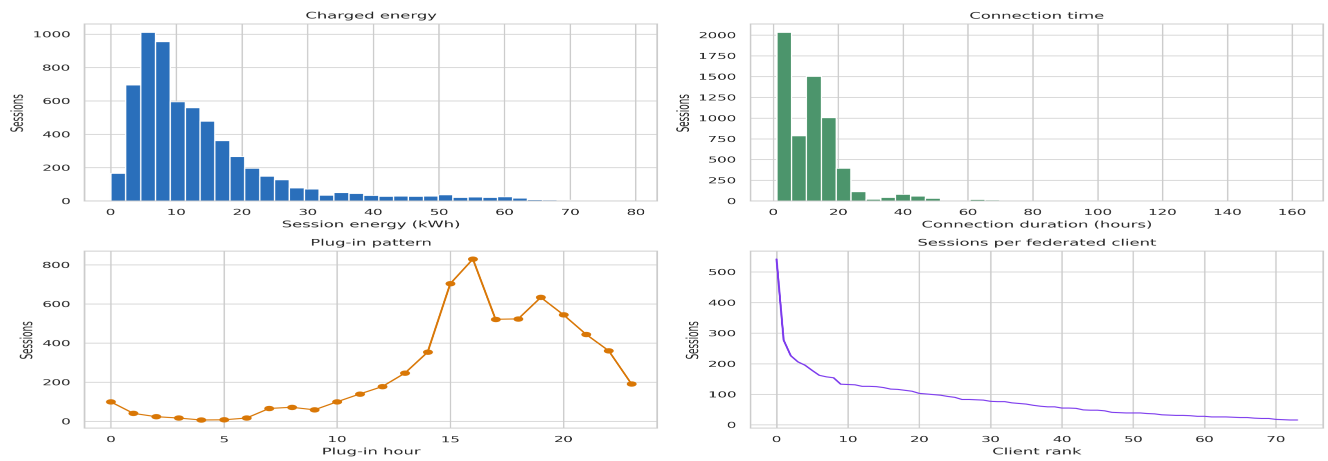

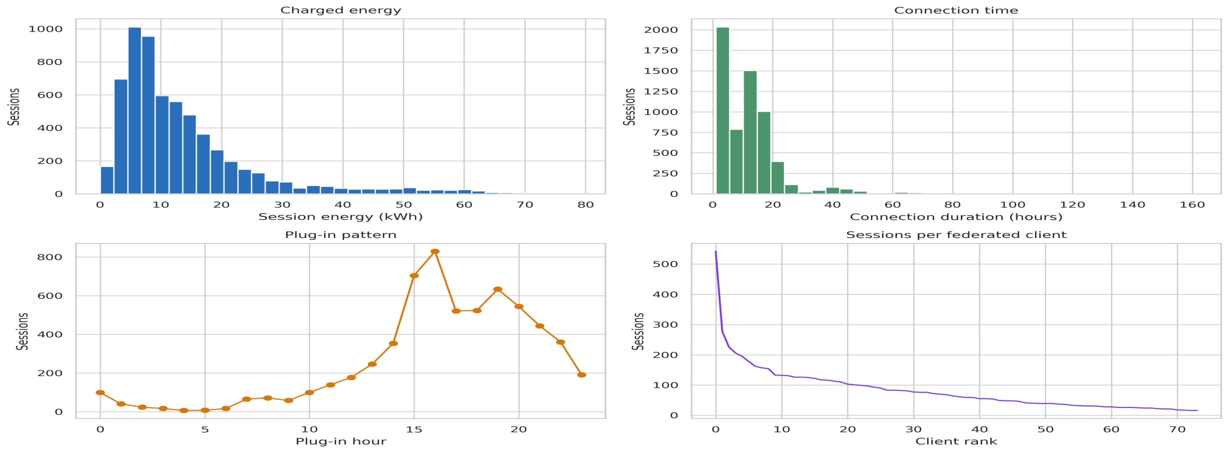

Figure 2 Measured charging-session distributions and client imbalance.

Figure 2 shows energy, duration, plug-in time, and client record counts. The long client-rank tail confirms uneven participation histories. Such imbalance can reduce the quality of a simple mean update. The plug-in profile also creates common evening demand. Long connection times still offer room for delayed charging. The scheduler uses that room without changing each recorded departure.

4.2 Residential background demand

Measured household imports come from Open Power System Data. The selected release contains six residential grid-import series. These meters report cumulative imported energy rather than hourly demand. The corrected workflow differences consecutive valid hourly readings. Negative changes and irregular time gaps are removed. Extreme increments receive a conservative data-quality filter. One hourly kilowatt-hour then equals one kilowatt of average demand.

A calendar profile uses median demand by month, weekday, and hour. The profile is scaled to a feeder serving 100 homes. This background load is paired with Norwegian charging sessions for simulation. The two datasets are not geographically co-located. Their combination forms a transparent composite feeder. Wiese et al. (2019) describe the open electricity data platform. The composite design supports reproducibility, but not local network validation.

4.3 Chronological splitting and preprocessing

Each client is split in time order. The first 70 percent of sessions form local training data. The next 15 percent form validation data. The final 15 percent form the test data. This process yields 4,281 training sessions and 924 validation sessions. The test set contains 958 later sessions. Chronological splits avoid the optimistic leakage caused by random session mixing.

Local history features are computed in a prequential manner. Each value uses only earlier records from the same client. The method never averages future test energy into an earlier feature. Energy targets receive a logarithmic transform before normalization. The transformation reduces the influence of rare high-energy sessions. A fixed seed controls sampling, initialization, and added noise. The seed value is 20260816.

Table 2: Dataset and feeder summary

Item	Value	Definition
Raw charging rows	6,878	Public charging reports
Clean charging sessions	6,163	After quality filters
Federated clients	74	At least 15 sessions each
Training sessions	4,281	First 70% per client
Validation sessions	924	Next 15% per client
Test sessions	958	Final 15% per client
Modeled homes	100	Scaled OPSD profile
Feeder capacity	136.06 kW	85% of uncontrolled peak

4.4 Training and scheduling configuration

The main experiment uses 40 federated rounds and two local epochs. Thirty percent of clients join each round. The update norm is clipped at 0.50. Private runs use noise multipliers of 0.5, 1.0, and 1.5. The main scheduling result uses the middle value. The feeder capacity equals 85 percent of the uncontrolled combined peak. Every vehicle uses a maximum charging power of 7.2 kW.

Table 3: Main experimental configuration

Parameter	Value
Centralized epochs	100
Federated rounds	40
Clients per round	23
Local epochs	2
Local learning rate	0.025
Update clipping norm	0.50
Main noise multiplier	1.0
Privacy delta	0.00001
Charger power	7.2 kW
Capacity ratio	0.85

4.5 Comparison cases

Five charging cases support the feeder comparison. Uncontrolled charging begins at arrival and continues until the recorded energy is delivered. Centralized smart charging uses predictions from pooled training. Federated smart charging uses nonprivate FedAvg predictions. Private federated charging uses clipped, noisy, securely aggregated updates. Oracle smart charging uses the recorded future energy as perfect knowledge. The oracle isolates scheduling potential from prediction error.

The centralized case provides an accuracy ceiling for the chosen network. The ordinary federated case isolates the cost of data separation. The private case then isolates the additional effect of clipping and noise. All predictive schedulers add a validation-based error margin. That margin is capped by the available connection time. The comparison uses identical test sessions and background demand.

5. RESULTS

5.1 Federated training behavior

Federated validation error changes quickly during the first training rounds. Later rounds show smaller and less stable improvements. The three private curves remain close to ordinary FedAvg. This pattern suggests that client differences dominate the added privacy noise. The result is consistent with known non-identical federated data problems (Li et al., 2020). More rounds alone may not close the centralized accuracy gap. Personalized heads or proximal objectives deserve further testing.

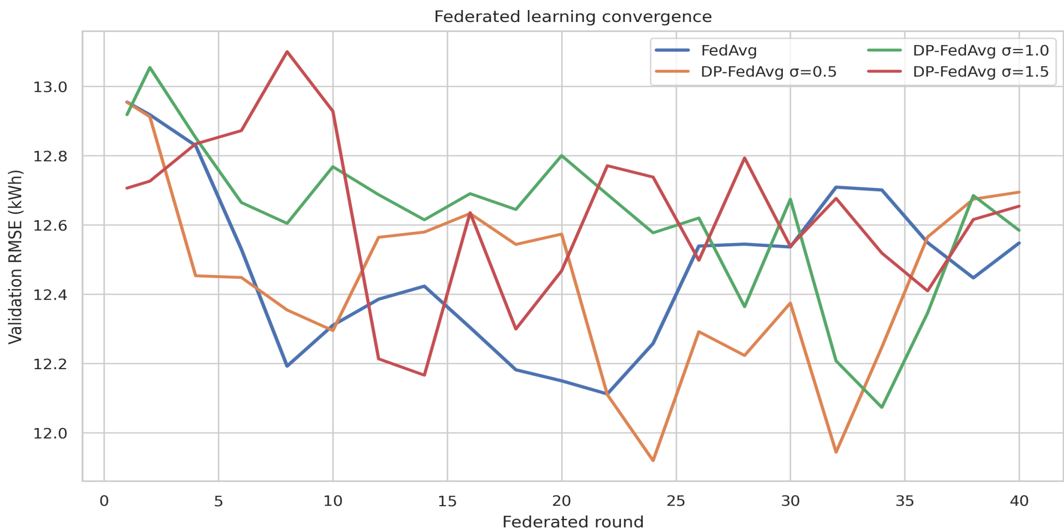

Figure 3 Validation error across federated training rounds

Figure 3 also shows that noise does not cause uncontrolled divergence. Every curve remains within a narrow error band after early training. The noise multiplier of 1.5 sometimes follows the nonprivate curve closely. This outcome does not mean stronger noise always improves accuracy. It reflects random optimization paths within one fixed-seed experiment. Repeated seeds are needed before comparing small curve differences.

5.2 Prediction accuracy

The centralized predictor performs best on all reported prediction errors. Its test MAE is 4.97 kWh, and its RMSE is 8.60 kWh. Its R squared value reaches 0.514. Ordinary FedAvg records 8.12 kWh MAE and 12.79 kWh RMSE. Its R squared value is slightly negative. The private model records 8.15 kWh MAE and 12.80 kWh RMSE. Its R squared value is also negative.

Table 4: Prediction results on 958 unseen sessions

Model	MAE kWh	RMSE kWh	R^2	P90 margin kWh
Centralized	4.97	8.60	0.514	10.46
FedAvg	8.12	12.79	−0.073	18.69
DP-FedAvg, σ=1.0	8.15	12.80	−0.075	18.35

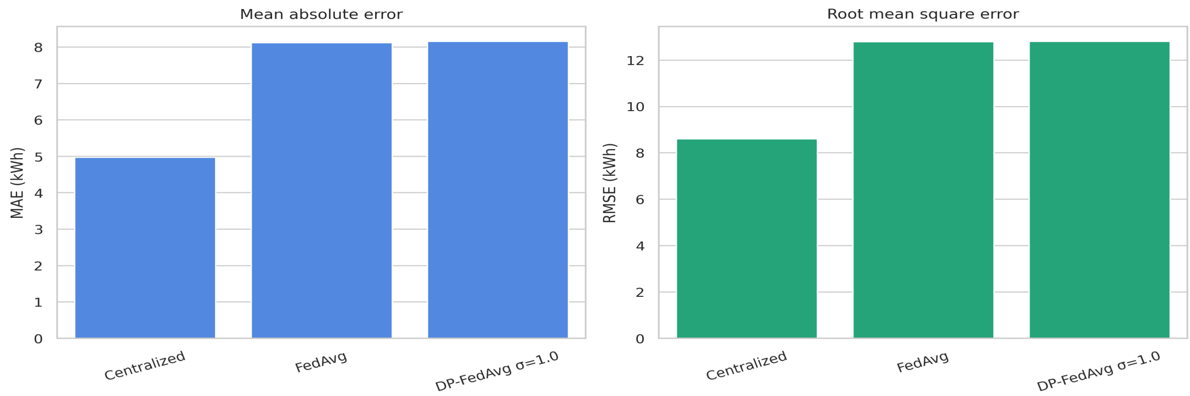

Figure 4 Final test errors for centralized and federated predictors.

The privacy mechanism adds almost no error beyond ordinary federated training. Private RMSE is only 0.02 kWh higher than FedAvg RMSE. The larger gap appears between centralized and federated models. Client histories contain unequal sizes and different energy patterns. Two local epochs can move updates toward these personal patterns. Simple averaging then produces a weak population estimate. The negative R squared values confirm that this baseline needs improvement.

5.3 Feeder load profile

Uncontrolled charging creates a combined peak of 160.07 kW. The modeled feeder capacity is 136.06 kW. Ten hourly intervals exceed that limit. All four smart cases keep hourly demand below the limit. The private schedule reaches a 100.55 kW peak. This level is 59.52 kW below the uncontrolled case. The resulting peak reduction is 37.18 percent.

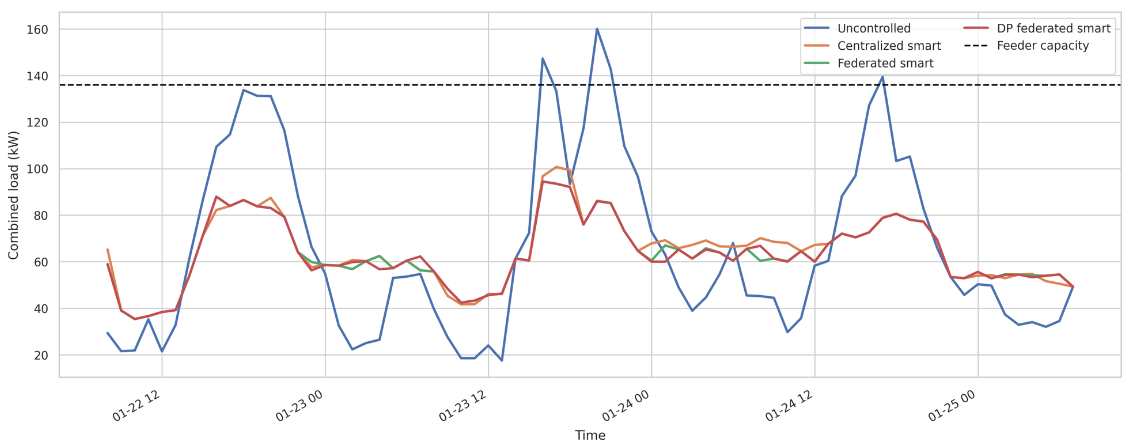

Figure 5 Seventy-two-hour feeder profile near the uncontrolled peak.

Figure 5 shows how the smart cases fill lower-load periods. Their lines remain below the dashed capacity limit. The three predictive schedules are visually close near the original peak. This similarity occurs despite their different prediction errors.

A conservative calibration margin reserves additional energy. The scheduler also uses long connection windows to avoid crowded slots. These controls can suppress peaks even when energy forecasts remain imperfect.

5.4 Congestion and load-shape measures

Centralized smart charging lowers the peak by 36.70 percent. Ordinary federated smart charging lowers it by 36.80 percent. Private federated smart charging achieves the largest measured reduction at 37.18 percent. This small ordering should not be treated as an accuracy ranking. Underprediction can reduce the scheduled peak while leaving energy unmet. Peak reduction must therefore be read beside service metrics.

Table 5: Grid and service performance

Method	Peak kW	Reduction %	PAR	Overload h	Unmet %	Delay h
Uncontrolled	160.07	0.00	4.96	10	0.05	0.41
Centralized smart	101.33	36.70	3.16	0	7.52	8.34
Federated smart	101.17	36.80	3.16	0	10.32	8.38
DP federated smart	100.55	37.18	3.14	0	10.57	8.37
Oracle smart	101.33	36.70	3.14	0	0.00	8.36

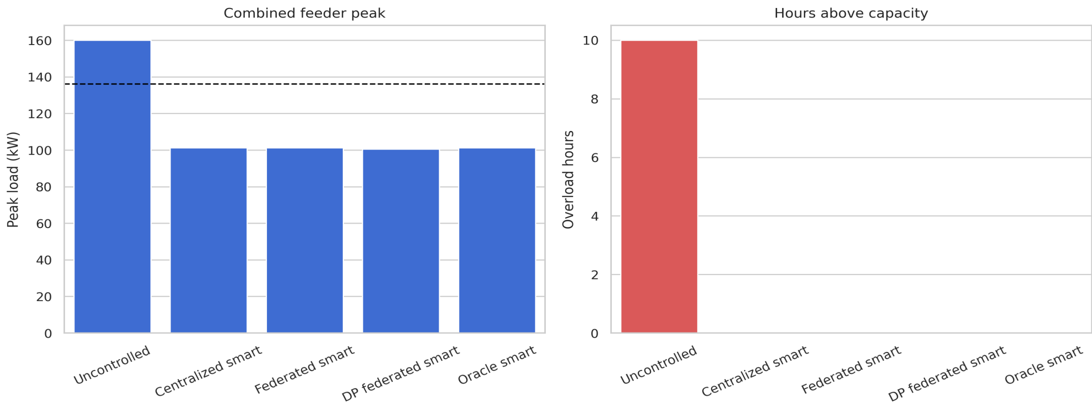

Figure 6: Combined feeder peaks and overload hours.

Private charging also lowers the peak-to-average ratio from 4.96 to 3.14. Its load variance falls from 504.86 to 366.16 square kilowatts. These changes show a broader load-shape improvement. They are not limited to one maximum hour. The flat capacity test produced the same outcome across tested ratios. The controlled peak remained below even the lowest tested capacity. A wider or tighter range is needed for an informative sensitivity curve.

5.5 Driver service effects

The oracle schedule delivers every requested kilowatt-hour. It proves that recorded connection windows contain enough physical flexibility. Predictive schedules fail to deliver some energy because requests remain imperfect. The centralized schedule leaves 7.52 percent unmet. The ordinary federated schedule leaves 10.32 percent unmet. The private schedule leaves 10.57 percent unmet. This final value equals about 1,453.5 kWh across the test horizon.

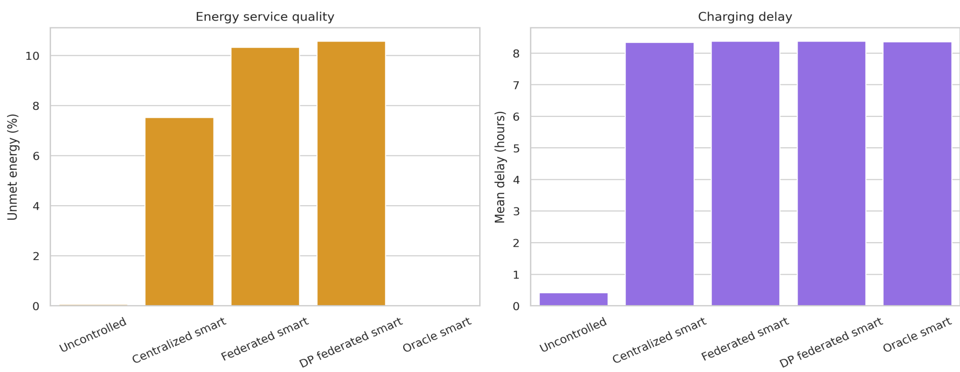

Figure 7 Unmet energy and mean charging delay by control method.

Every smart schedule adds roughly eight hours of energy-weighted delay. That delay reflects movement toward quiet overnight periods. Long connections make such movement technically possible. Drivers may still reject the result when urgent departures occur. A practical controller should include minimum state-of-charge targets. It should also protect a near-term energy reserve. Those constraints would reduce flexibility but improve acceptance.

5.6 Privacy and utility

The weakest tested noise setting gives an epsilon value of 52.41. The middle setting reduces epsilon to 14.91. The strongest setting reduces epsilon further to 7.92. All values use a delta of 0.00001. Test RMSE remains between 12.80 and 12.93 kWh. The narrow range shows a mild measured utility cost. However, every privacy budget remains conditional on the stated training design.

Table 6: Differential privacy and prediction utility

Noise σ	Epsilon	Delta	MAE kWh	RMSE kWh	Mean clip factor
0.5	52.41	0.00001	8.07	12.93	0.352
1.0	14.91	0.00001	8.15	12.80	0.345
1.5	7.92	0.00001	8.16	12.90	0.334

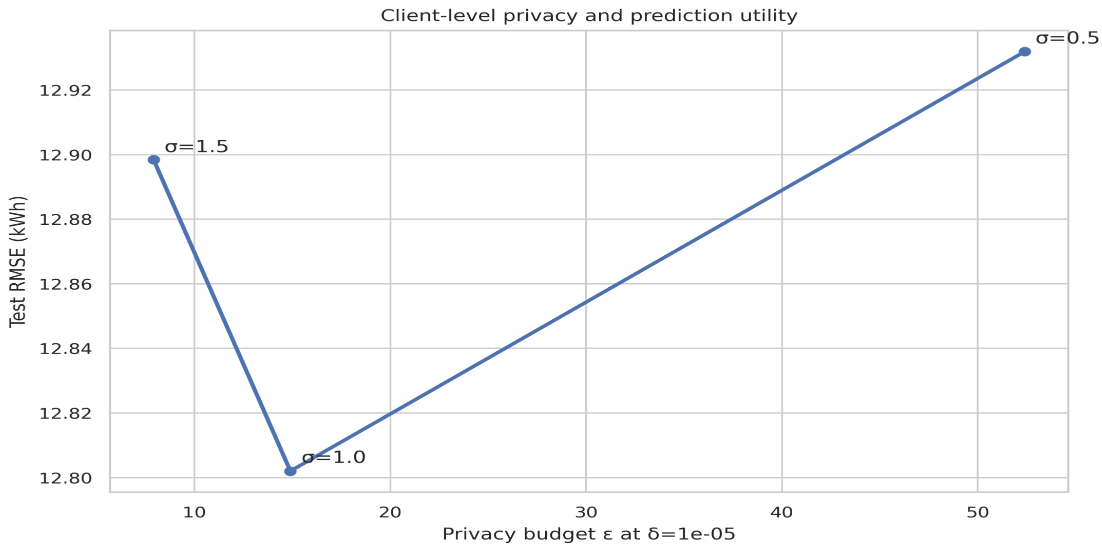

Figure 8 Client-level privacy budget and prediction error.

The noise multiplier of 1.5 offers the strongest tested privacy. Its RMSE is only 0.10 kWh above the middle setting. That result makes it an attractive candidate for further trials. Still, one fixed seed cannot establish a stable privacy ranking. Multiple seeds should report mean errors and confidence intervals. Client participation should also vary across realistic availability patterns.

5.7 Secure aggregation and load duration

Pairwise mask cancellation remains numerically accurate in every run. Ordinary FedAvg shows a maximum difference of 0.000000119. Private runs show errors near one quadrillionth. These values confirm correct cancellation within floating-point arithmetic. They do not demonstrate resistance to active attacks. Production security needs authenticated key exchange and dropout recovery. It also needs formal implementation review.

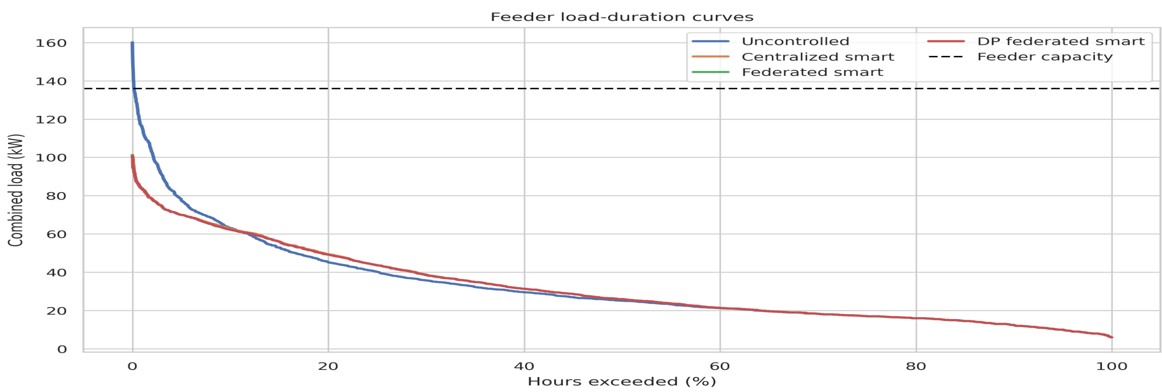

Figure 9 Feeder load-duration curves across the test horizon.

The load-duration curves show how often each load level is exceeded. Uncontrolled charging crosses the capacity line during a short upper tail. Every smart case stays below that line. The curves then cluster across most remaining hours. This result confirms that control mainly reshapes the most congested periods. It also avoids claiming a uniform demand reduction. Charging energy is shifted or missed, rather than freely removed.

6. DISCUSSION

6.1 Main finding

The proposed method removes all ten modeled overload hours. It also reduces the combined feeder peak by 37.18 percent. These values answer the main grid question positively. Private federated predictions can support useful feeder control. Raw session records do not need central storage for that control. However, the service result prevents an unconditional success claim. More than one tenth of requested energy remains unmet.

The oracle case explains this outcome. Perfect energy knowledge reaches a similar peak while delivering all requested energy. Therefore, feeder capacity is not the main service barrier in this test. Prediction error creates most missed energy. The calibration margin reduces underestimation but cannot remove it. Improving the federated predictor should therefore improve both fairness and service. Peak performance alone cannot reveal that need.

6.2 Centralized advantage and federated heterogeneity

Centralized training has direct access to all 4,281 training sessions. It can learn broad patterns from large and small clients together. FedAvg instead performs several steps on separate personal histories. Unequal sample counts and different behavior move those updates in conflicting directions. Clipping further limits the strongest updates. The resulting global model has much higher error than centralized training. This gap matches the heterogeneity problem described by Li et al. (2020).

Several improvements could address that gap. FedProx could restrict local drift through a proximal term. Personalized output layers could retain user-specific energy patterns. Clustered federation could group clients with similar charging behavior. Robust aggregation could reduce the influence of unstable updates. A recurrent model might also capture longer personal sequences. Each option should retain the same chronological test design.

6.3 Meaning of the privacy result

Federated learning keeps raw examples at their source. Secure aggregation hides individual protected updates from the server. Differential privacy bounds one client's effect on the released model. These protections address related but different disclosure paths. The reported epsilon values make the privacy claim measurable. They do not make the system anonymous or risk-free. Identity, authentication, device logs, and scheduling commands still need protection.

The strongest tested noise setting deserves more attention. It lowers epsilon from 14.91 to 7.92 with little measured error change. That observation supports further repeated trials at higher noise. It does not prove that epsilon 7.92 is acceptable for every utility. Privacy targets should follow legal duties and a stated threat model. Participation frequency should also be included in long-term accounting.

6.4 Practical implications

A practical deployment could place clients inside chargers, vehicles, or home energy gateways. The feeder operator would publish a model and capacity signal. Clients would return protected updates during scheduled training windows. A local controller could then apply the global predictor. Drivers should provide minimum energy and departure preferences. Those preferences should become hard service constraints. The present lowest-load rule could then operate on remaining flexibility.

Utilities should also avoid using one static capacity threshold. Transformer temperature, phase loading, solar output, and faults change safe limits. A field controller should receive updated operating envelopes. It should fall back safely when communication fails. Local emergency charging must remain possible. Audit logs should record decisions without storing unnecessary travel history. These safeguards affect acceptance as much as model accuracy.

6.5 Limitations

The experiment is a composite simulation, not a field trial. Norwegian charging sessions are paired with German household demand. The feeder contains 100 modeled homes and one aggregate capacity. It does not model voltage, reactive power, phases, or thermal state. Network losses and charger efficiency are also omitted. The results therefore describe active-power congestion under stated assumptions. They should not be treated as a complete distribution study.

The secure aggregation component is an arithmetic protocol emulation. It does not use production cryptography or test client dropout. The privacy accountant assumes the stated sampling and Gaussian mechanism. Only one fixed seed supports the reported main results. No confidence intervals or statistical tests are available. The scheduler also ignores prices and driver compensation. A deployment study must address each limitation before operational claims.

The capacity sensitivity range did not bind the controlled schedule. All tested ratios produced the same private control metrics. That flat result reflects ample headroom after peak shifting. It does not establish robustness under severe constraints. Future experiments should test lower limits and time-varying envelopes. They should also include higher charging powers and dense adoption.

7. CONCLUSION

This paper examined private federated learning for residential electric vehicle charging. The method kept raw charging sessions on their clients. It combined update clipping, distributed Gaussian noise, and secure-sum masking. A feeder-aware scheduler then moved predicted demand into lower-load hours. The evaluation used 6,163 measured sessions and open household imports. Correct conversion of cumulative meters produced a plausible 160.07 kW uncontrolled peak.

Private smart charging lowered that peak to 100.55 kW. It removed ten overload hours and reduced the peak by 37.18 percent. The main privacy setting achieved epsilon 14.91 at delta 0.00001. Its prediction RMSE was 12.80 kWh. Yet the schedule left 10.57 percent of requested energy unmet. This service loss is the central caution from the experiment.

Future work should improve learning across unequal client histories. It should test personalized or proximal federated models across several random seeds. Charging control should include minimum energy guarantees and urgent departure protection. Field trials should use real feeder topology, transformer temperature, and secure communication. With those additions, private coordination could reduce congestion without exposing daily travel behavior.

Declarations

Data availability

The electric vehicle data are publicly available from Mendeley Data. The dataset identifier is 10.17632/jbks2rcwyj.3. Household imports are available from Open Power System Data. The release date is April 15, 2020. The supplied notebook and result archive contain the full experiment outputs.

Code and reproducibility

The notebook records every setting and uses seed 20260816. It downloads public data, trains all baselines, and exports ten figures. It also exports six result tables and one configuration file. The reported paper values come from experiment version 2.0.

Conflict of interest

The author declares no conflict of interest.

Funding

No external funding was used for this experiment.

REFERENCES

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 308–318. https://doi.org/10.1145/2976749.2978318
2. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175–1191. https://doi.org/10.1145/3133956.3133982
3. Brinkel, N. B. G., Schram, W. L., AlSkaif, T. A., Lampropoulos, I., & van Sark, W. G. J. H. M. (2020). Should we reinforce the grid? Cost and emission optimization of electric vehicle charging under different transformer limits. Applied Energy, 276, 115285. https://doi.org/10.1016/j.apenergy.2020.115285
4. Clement-Nyns, K., Haesen, E., & Driesen, J. (2010). The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. IEEE Transactions on Power Systems, 25(1), 371–380. https://doi.org/10.1109/TPWRS.2009.2036481
5. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407. https://doi.org/10.1561/0400000042
6. Gan, L., Topcu, U., & Low, S. H. (2013). Optimal decentralized protocol for electric vehicle charging. IEEE Transactions on Power Systems, 28(2), 940–951. https://doi.org/10.1109/TPWRS.2012.2210288
7. Hilshey, A. D., Hines, P. D. H., Rezaei, P., & Dowds, J. R. (2013). Estimating the impact of electric vehicle smart charging on distribution transformer aging. IEEE Transactions on Smart Grid, 4(2), 905–913. https://doi.org/10.1109/TSG.2012.2217385
8. Abdussalam Ali Ahmed (2025). From Transition to Transformation: A Comparative Engineering Study of Hybrid and Electric Vehicles. Libyan Open University Journal of Applied Sciences (LOUJAS), 1(1), 11-19. https://doi.org/10.65422/loujas.v1i1.48
9. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083
10. Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of Machine Learning and Systems, 2, 429–450. https://proceedings.mlsys.org/paper/2020/hash/1f5fe83998a09396ebe6477d9475ba0c-Abstract.html
11. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-

efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 54, 1273–1282. https://proceedings.mlr.press/v54/mcmahan17a.html
12. Mironov, I. (2017). Rényi differential privacy. 2017 IEEE 30th Computer Security Foundations Symposium, 263–275. https://doi.org/10.1109/CSF.2017.11
13. Muratori, M. (2018). Impact of uncoordinated plug-in electric vehicle charging on residential power demand. Nature Energy, 3, 193–201. https://doi.org/10.1038/s41560-017-0074-z
14. Richardson, D. B. (2013). Electric vehicles and the electric grid. Renewable and Sustainable Energy Reviews, 19, 247–254. https://doi.org/10.1016/j.rser.2012.11.042
15. Azar Ali Matoug, & Abdussalam Ali Ahmed. (2026). Simulation-Based Performance Analysis of Phosphoric Acid Fuel Cells: Polarization Behavior, Power Density, and Sensitivity to Operating Conditions. African Union Journal of Academic and Research Studies, 1(2), 1-8.
16. Sadeghian, O., Oshnoei, A., Mohammadi-Ivatloo, B., Vahidinasab, V., & Anvari-Moghaddam, A. (2022). A comprehensive review on electric vehicles smart charging. Journal of Energy Storage, 54, 105241. https://doi.org/10.1016/j.est.2022.105241
17. Sortomme, E., & El-Sharkawi, M. A. (2011). Optimal charging strategies for unidirectional vehicle-to-grid. IEEE Transactions on Smart Grid, 2(1), 131–138. https://doi.org/10.1109/TSG.2010.2090910
18. Sørensen, Å. L., Lindberg, K. B., Sartori, I., & Andresen, I. (2021a). Residential electric vehicle charging datasets from apartment buildings. Data in Brief, 36, 107105. https://doi.org/10.1016/j.dib.2021.107105
19. Sørensen, Å. L., Lindberg, K. B., Sartori, I., & Andresen, I. (2021b). Analysis of residential electric vehicle energy flexibility potential based on real-world charging reports and smart meter data. Energy and Buildings, 241, 110923. https://doi.org/10.1016/j.enbuild.2021.110923
20. Wiese, F., Schlecht, I., Bunke, W.-D., Gerbaulet, C., Hirth, L., Jahn, M., Kunz, F., Lorenz, C., Mühlenpfordt, J., Reimann, J., & Scharf, M. (2019). Open Power System Data: Frictionless data for electricity system modelling. Applied Energy, 236, 401–409. https://doi.org/10.1016/j.apenergy.2018.11.097
21. Abdulgader Alsharif, Abdussalam Ali Ahmed, Omar Ahmed Mohamed, & Taha Muftah Abuali. (2026). Hybrid Machine Learning Approaches for Accurate Solar Energy Forecasting from Real-World Weather Data. Libyan Journal of Health, Science, and Development (LJHSD), 2(1), 09-17