POLITECNICO DI TORINO
Repository ISTITUZIONALE

Privacy-Preserving Federated Learning for Household Characteristic Identification

Original
Privacy-Preserving Federated Learning for Household Characteristic Identification / Malan, E., De Vizia, C., Castangia, M., Peluso, V., Calimera, A., Macii, E.. - (2025), pp. 2040-2045. (2025 IEEE 49th Annual Computers, Software, and Applications Conference (COMPSAC) Toronto ON (CAN) 08-11 July 2025) [10.1109/compsac65507.2025.00285].

Availability:
This version is available at: 11583/3002806 since: 2025-09-04T15:00:29Z

Publisher:
IEEE

Published
DOI:10.1109/compsac65507.2025.00285

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
IEEE postprint/Author's Accepted Manuscript

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

18 November 2026

Privacy-Preserving Federated Learning for Household Characteristic Identification

Erich Malan, Claudia De Vizia, Marco Castangia*, Valentino Peluso, Andrea Calimera, Enrico Macii*
Department of Control and Computer Engineering, Politecnico di Torino, Turin, Italy
*Interuniversity Department of Regional and Urban Studies and Planning, Politecnico di Torino, Turin, Italy
{erich.malan, claudia.devizia, marco.castangia, valentino.peluso, andrea.calimera, enrico.macii}@polito.it

***Abstract*—This work presents a privacy-preserving training framework for household characteristic identification from electricity consumption data. The proposed framework integrates two main components: (i) a synthetic data generation pipeline capable of replicating realistic energy traces from diverse family compositions, capturing fine-grained sociodemographic attributes such as household size, employment status, age groups, and home occupancy patterns; (ii) a training strategy based on Federated Learning (FL) secured with homomorphic encryption, enabling collaborative model training while preserving data ownership. Our synthetic dataset enables the performance assessment of different training scenarios, including siloed model training by individual energy utilities and secure collaboration via FL. Experimental results show that siloed training leads to inconsistent and suboptimal performance, while privacy-preserving FL achieves accuracy comparable to conventional centralized training—an ideal yet not viable option due to data regulation constraints. Our findings highlight the effectiveness of FL as a secure solution for collaborative sociodemographic profiling in smart grids.**

***Index Terms*—Household Characteristic Identification, Smart Grids, Federated Learning, Homomorphic Encryption**

I. Introduction

The energy consumption profiles of residential buildings are closely related to the household characteristics. In fact, sociodemographic factors such as employment status, income level, and age significantly influence how energy is used throughout the day. However, this information is rarely available, as it typically relies on costly surveys that often face consumers resistance. Nevertheless, these household attributes can be inferred with reasonable accuracy from smart meter load profiles [1]. For instance, a household comprising two full-time workers would generally exhibit lower energy usage during midday hours, whereas a retired couple is likely to display higher daytime consumption. These behavioral patterns offer utilities the opportunity to infer customer characteristics directly from smart meter data.

Energy utilities and retailers have a vested interest in identifying customer profiles for various strategic purposes. Access to detailed household information enables the design of more effective demand-response programs tailored to specific consumer segments [2]. Such customization can reduce the disruption of load shifting on users' daily routines. Furthermore, retailers can offer personalized pricing structures that better reflect the economic capacity of different customer segments. Tailored incentives and service plans can also enhance customer retention by aligning more closely with individual expectations and needs. Additionally, retailers can identify atypical consumption patterns within a segment, rewarding energy-conscious consumers and deliver targeted guidance to those exhibiting excessive usage [3]. Importantly, consumers themselves benefit from this profiling through more responsive and efficient service offerings.

In deregulated energy markets, consumers have the autonomy to select their energy provider. As a consequence, retailers compete by offering more attractive tariffs and services. While this promotes personalized offers, lower prices, and a broader range of choices for consumers, it also results into fragmented data collection, as smart meter data is managed by a variety of local energy utilities. Privacy regulations and confidentiality concerns often prevent utilities from sharing customer data with external entities, thereby limiting the possibility to analyze and infer household characteristics beyond their own customer base [4].

Federated Learning (FL) represents a promising solution to this challenge by enabling collaborative model training across multiple parties with no need to share sensitive raw data [5]. This privacy-preserving approach allows utilities to develop more accurate and generalizable models of household behavior by leveraging broader and more diverse datasets. As a result, it becomes possible to enhance service personalization and improve energy system efficiency while maintaining compliance with data privacy standards.

Consequently, this work explores the potential of FL to train machine-learning services for household profile identification. Our contributions are twofold. First, we introduce a synthetic data generation pipeline capable of modeling realistic electricity consumption traces. The synthetic dataset is generated using an agent-based simulator [6], [7], which characterizes each household member by sociodemographic attributes and accurately reproduces consumption patterns at a fine level of granularity. This dataset enables accurate characterization of consumption behaviors and serves as a benchmark for evaluating collaborative training among energy utilities. Second, we implement a privacy-preserving FL framework secured with Homomorphic Encryption (HE), enabling distributed model training without exposing sensitive customer data.

Experimental results show that models trained in isolation by individual utilities suffer from inconsistent and low accuracy due to data fragmentation and heterogeneity. In contrast, privacy-preserving FL reaches performance levels close to an

ideal centralized training setup, proving its potential for household characteristic identification in distributed energy systems.

The paper is structured as follows. Section II provides background information and a review of related work. Section III introduces the framework, describing the synthetic dataset and the proposed FL-based method. Section IV presents the experimental setup and the collected results. Finally, Section V draws conclusions.

II. Background & Related Work

A. Machine Learning for Electricity Consumption Analysis

Previous works have addressed the problem of identifying household characteristics by classifying daily power profiles from smart meter data. Classification methods offer a detailed characterization of consumer profiles by identifying key sociodemographic attributes of household members. Such methods can infer valuable information, including household size, employment status, age distribution, and income level. This enriched understanding enables utility providers to perform more precise household profiling and develop more tailored and effective energy solutions for the end users. Pekey et al. [8], for instance, evaluated multiple classification algorithms for predicting key household attributes using smart meter readings. Their approach involved aggregating raw load measurements by computing statistical features (such as the average, minimum, and maximum) across various temporal windows throughout the day. Similar feature engineering strategies have been adopted in other studies, often in conjunction with support vector machine classifiers [9]. Conversely, deep learning methods are particularly well-suited for automatic extraction of relevant features during training. Deep learning methods allow for direct processing of raw load data, thus bypassing manual feature engineering, establishing as the standard processing favored by prior work. Within this context, one-Dimensional Convolutional Neural Networks (1D-CNNs) have emerged as the most common architecture for household characteristic identification [10], [11]. Notably, CNNs excel at constructing hierarchical representations from raw time series data, making them particularly effective for analyzing large volumes of power consumption profiles. More recently, CNNs have also been integrated with long short-term memory (LSTM) networks to enhance the modeling of temporal dynamics in the data [12].

B. Federated Learning

FL enables multiple organizations to collaborate in training a global machine-learning model without sharing proprietary data. Each organization trains a local instance of the model using its private dataset and sends model updates to a central server, which aggregates the received updates to gradually refine the global model. This process is iterative and involves several synchronization rounds between the server and the participating parties. A differentiating feature of FL, in contrast to traditional centralized training, is that it relies on the aggregation of locally trained models rather than the centralization of distributed datasets. This feature preserves

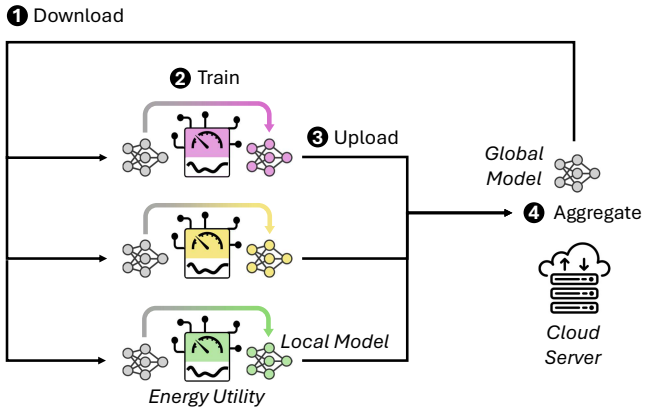


Fig. 1. Federated Learning.

data locality and enhances privacy, facilitating the deployment of machine-learning services across organizational boundaries in compliance with data protection regulations.

Algorithm 1 presents the pseudo-code for a baseline implementation of synchronous FL, following the commonly used Federated Averaging (FedAvg) strategy [13]. As schematically illustrated in Fig. 1, the workflow involves two main entities: a central server that coordinates model aggregation and synchronization, and a set of participating parties $\mathcal{P}$, each owning a private dataset $\mathcal{D}_p$ used for model training. In our work, the participating parties are energy utility companies owning datasets of household electricity consumption data recorded by smart meters. Initially, the server assigns random values to the global model weights $\boldsymbol{\omega}^1$ (line 1). The training process is then organized into a sequence of R synchronization rounds (lines 2–9). At the beginning of each round r, the server distributes the current global model $\boldsymbol{\omega}^r$ to the parties, who locally store it as $\boldsymbol{\omega}_p^{r,1}$ (line 4). Each participant proceeds to perform S local training steps on its dataset $\mathcal{D}_p$, generating an updated local model $\boldsymbol{\omega}_p^{r,S}$ (line 5). After training, each participant computes a local model update $\boldsymbol{\Delta}_p^r$ as the difference between the updated and the received model weights (line 6), and uploads $\boldsymbol{\Delta}_p^r$ to the server (line 7). After collecting updates from all the parties, the server aggregates them through averaging according to the FedAvg protocol. The resulting global update $\boldsymbol{\Delta}^r$, referred to as the *pseudo-gradient*, is then used to refine the global model generating the updated weights $\boldsymbol{\omega}^{r+1}$ for the next round (lines 8–9).

FL has recently gained traction in the smart grid domain, finding applications across a wide spectrum of use cases, including anomaly detection, load forecasting, and household profiling.

In the context of anomaly detection, [14] proposed an unsupervised FL framework to identify irregular patterns from data collected by remote terminal units in substations. Another study [15] focused on detecting cyberattacks in electrical transmission systems, while [16] addressed the problem of energy theft detection from electricity consumption data.

Concerning load forecasting, FL have been exploited for training on distributed smart meter data to predict household-

Algorithm 1: Federated Learning with FedAvg.

/* Server initialization */
$\boldsymbol{\omega}^1 \leftarrow$ Random model weights
/* Synchronization Rounds */
for $r = 1,...,R$ **do**
  /* Local Optimization */
  for $p \in \mathcal{P}$ **do in parallel**
    $\boldsymbol{\omega}_p^{r,1} \leftarrow \boldsymbol{\omega}^r$
    $\boldsymbol{\omega}_p^{r,S} \leftarrow \text{TRAIN}(\boldsymbol{\omega}_p^{r,1}, \mathcal{D}_p, S)$
    $\boldsymbol{\Delta}_p^r \leftarrow \boldsymbol{\omega}_p^{r,S} - \boldsymbol{\omega}^{r-1}$
    Upload $\boldsymbol{\Delta}_p^r$ to the server
/* Server Aggregation */
$\boldsymbol{\Delta}^r \leftarrow \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \boldsymbol{\Delta}_p^r$
$\boldsymbol{\omega}^{r+1} \leftarrow \boldsymbol{\omega}^r + \boldsymbol{\Delta}^r$

level energy demand [17] and to predict net energy consumption in residential solar photovoltaic systems [18] for helping utilities in anticipating local production and demand.

Research in household profiling has leveraged FL to analyze behavioral and structural characteristics of households. The authors of [19] proposed an unsupervised approach for clustering customers with similar consumption patterns. Other works have targeted the inference of dwelling properties, appliance efficiency, or high-level sociodemographic attributes from smart meter traces [20]. Our work differs as it aims to infer fine-grained information about household compositions.

C. Homomorphic Encryption

The data locality promoted by FL alone does not ensure full privacy protection. Adversaries such as honest-but-curious servers can perform attacks like gradient inversion [21], which may extract sensitive information from the received model updates. Therefore, additional defensive measures are required to strengthen privacy protection. Homomorphic Encryption (HE) is one such technique that has gained attention for its potential to mitigate these risks, as it prevents access to raw updates on the server side.

HE is a cryptographic technique that allows computations on encrypted data without prior decryption. Among the existing schemes, the Paillier [22] and Cheon-Kim-Kim-Song (CKKS) [23] algorithms emerged as the most widely adopted in the context of FL.

The deployment of HE relies on a trusted authority to generate and manage a pair of cryptographic keys—public and private—that are used during the encryption, decryption, and evaluation phases. In the encryption phase, plaintext data is transformed into ciphertext using the public key. Decryption, which can only be performed by entities in possession of the private key, restores the data to its original form. The evaluation phase, typically carried out by untrusted entities, such as the FL server, enables operations on ciphertexts without exposing the underlying data. The trusted authority manages the generation, secure storage, and distribution of the keys among the participating parties.

Integrating HE into FL frameworks raises several concerns as cryptographic operations introduce huge communication and computation overheads. While this work does not focus on HE optimization, we summarize the main existing strategies in the following text. A primary source of inefficiency in HE lies in the large size of encrypted messages, leading up to $35\times$ more data traffic in federated settings [24]. An effective optimization to mitigate this problem is the technique of batching [25], which reduces the size of transmitted encrypted messages by packing multiple plaintext values into a single ciphertext. In our work, we used the CKKS scheme implemented by open-source TenSEAL library [26], which provides native support for batching. Complementary optimizations can be applied to further reduce HE overheads in FL, with techniques operating at different stages of the stack, such as efficient key generation [27], partial encryption [28], and selective synchronization [24], [29], [30].

III. Training Framework

A. Dataset Description

Detailed datasets that include both electricity consumption and household composition characteristics are often scarce and, when publicly available, may not be large enough to perform robust algorithm testing. Consequently, researchers often find themselves testing algorithms on the same limited datasets. Synthetic datasets present a viable alternative that allows researchers to generate arbitrarily large datasets to meet the testing requirements. Moreover, they provide full control over critical factors like class imbalance and data composition, enabling more comprehensive evaluations of algorithm performance.

For this reason, this work uses a dataset generated by an agent-based simulator presented in [6], [7], which is capable of replicating the electricity consumption patterns of the residential population starting from data usually available in almost all countries, i.e. the Time of Use diaries [31]. In these diaries, a sample population documents activities performed over 24 hours, such as sleeping and cooking, with a resolution of 10 minutes.

To replicate household electricity consumption patterns, the selected simulator includes two types of agents: i) the user agent, which replicates the activities of real individuals using a semi-Markov model [32], and ii) the household agent, which converts these activities into specific appliance consumption. The user agents are characterized by gender and employment, as both features have been proven to impact the time spent at home and the activities - thus, the electricity consumption [33]. For employment status, individuals can be classified as full-time, part-time, students, unemployed/retired, children, or homemakers.

For the study proposed in this paper, the Italian Time of Use diaries [34] were utilized to generate a dataset that reflects the electricity consumption of some Italian families. The resulting dataset comprises one year of electricity usage data of 1000 households, originally generated at 10-minute intervals and subsequently resampled to a hourly resolution.

According to *Istituto Nazionale di Statistica*, over the last twenty years, the proportion of single-person households has increased, while the percentage of larger households has

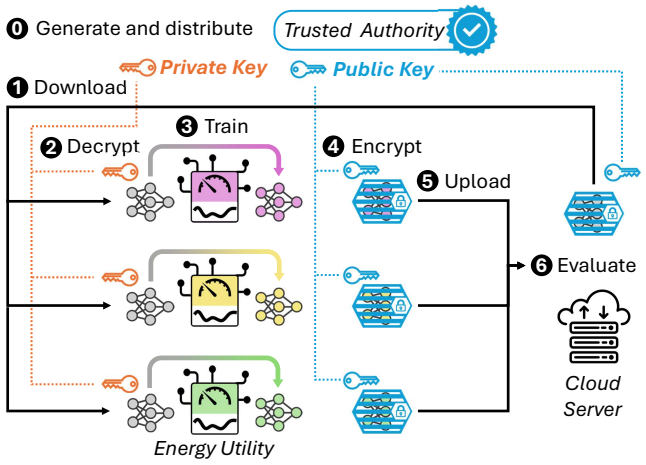


Fig. 2. Federated Learning with Homomorphic Encryption.

decreased. In fact, single-person households represent more than 36% of the total families, while households with five or more members represent less than 5% of the total [35]. Thus, it has been decided to include only families with up to four members in the dataset, for testing purposes. More specifically, we considered the following household compositions:

- Single-person households: full-time (10%), part-time (5%), student (5%) and retired (5%).
- Two-person households: a couple of retired individuals (5%), a couple of students (5%), a couple of full-time workers (5%), a couple with a full-time worker and a homemaker (5%), and a full-time worker with a child (5%).
- Three-member household: i) a couple of full-time workers with a child (10%), ii) a single parent with two children (5%), iii) a full-time worker, a partial worker, and a child (5%), iv) a full-time worker, a partial worker, and an unemployed/retired person (5%).
- Four-member household: i) two full-time workers and two children (5%), ii) a full-time worker, a partial worker, and two children (10%), iii) a full-time worker, a homemaker, a student, and a child (5%) and iv) two part-time workers, a student, and a child (5%).

While these combinations represent only a subset of all the possible ones, they have been chosen to cover all principal combinations in terms of user time availability and children presence. These factors are typically the key variables of interest, as they often influence decisions to adjust electricity consumption in response to requests [36] or variable tariffs [37].

B. Privacy-Preserving FL with HE

Our goal is to train a machine-learning classifier to infer the composition of a household from the electricity consumption profiles. To enable privacy-preserving collaboration among energy utilities, we rely on a FL architecture featuring HE protection. Fig. 2 illustrates its schematic diagram, while the pseudo-code in Algorithm 2 outlines the processing flow. Changes

Algorithm 2: Privacy-Preserving Federated Learning with Homomorphic Encryption.

1 /* Trusted Authority */
2 Generate private (sk) and public (pk) keys
3 Distribute sk and pk to the parties
4 Distribute pk to the server
/* Server initialization */
5 $\boldsymbol{\omega}^1 \leftarrow$ Random model weights
/* Synchronization Rounds */
6 **for** $r = 1,\dots,R$ **do**
/* Local Optimization */
7 **for** $p \in \mathcal{P}$ **do in parallel**
8 $\tilde{\boldsymbol{\omega}}_p^{r,1} \leftarrow \tilde{\boldsymbol{\omega}}^r$
9 $\boldsymbol{\omega}_p^{r,1} \leftarrow \textsc{Decrypt}(\tilde{\boldsymbol{\omega}}_p^{r,1}, sk)$
10 $\boldsymbol{\omega}_p^{r,S} \leftarrow \textsc{Train}(\boldsymbol{\omega}_p^{r,1}, \mathcal{D}_p, S)$
11 $\boldsymbol{\Delta}_p^r \leftarrow \frac{1}{|\mathcal{P}|}(\boldsymbol{\omega}_p^{r,S} - \boldsymbol{\omega}^r)$
12 $\tilde{\boldsymbol{\Delta}}^r \leftarrow \textsc{Encrypt}(\boldsymbol{\Delta}_p^r, pk)$
13 Upload $\tilde{\boldsymbol{\Delta}}_p^r$ to the server
/* Server Optimization */
14 $\tilde{\boldsymbol{\Delta}}^r \leftarrow \textsc{Evaluate}(\sum_{p\in\mathcal{P}} \tilde{\boldsymbol{\Delta}}_p^r, pk)$
15 $\tilde{\boldsymbol{\omega}}^{r+1} \leftarrow \textsc{Evaluate}(\tilde{\boldsymbol{\omega}}^r + \tilde{\boldsymbol{\Delta}}^r, pk)$

with respect to the baseline FL implementation (Algorithm 1) are highlighted with gray background. The process begins with a Trusted Authority responsible for generating the public and private key pair, along with the related cryptographic context that defines the configuration settings required for operations over encrypted data. The private key is used for decryption, while the public key enables encryption and computation over ciphertext. For brevity, only the key generation step is reported (line 2). The key authority distributes both the private key (sk) and public key (pk) to the training parties, while the server receives the public key only (lines 3–4).

At each round, the participating parties download the encrypted global model weights $\tilde{\boldsymbol{\omega}}^r$ from the server (line 8) and decrypt them using the private key (line 9), before retraining the model on the local data (line 10). After training, each party proceeds to the upload phase, which includes three main steps (lines 11–13). First, the model updates $\boldsymbol{\Delta}_p^r$ are rescaled locally, offloading this computational burden from the server. The updates are then encrypted, resulting in ciphertexts $\tilde{\boldsymbol{\Delta}}_p^r$ (line 16), which are subsequently sent to the server (line 13). Once all encrypted updates have been received, the server performs a secure aggregation (line 14) to compute the updated global model in encrypted form $\tilde{\boldsymbol{\omega}}^{r+1}$ (line 15).

IV. Experimental Setup & Results

A. Data Partitioning

The proposed synthetic dataset includes 17 family compositions representing the target classes for household characteristic identification. For training and evaluation, the dataset was split into training and test sets. Specifically, for each of the 17 classes, 5 households were randomly selected to form a balanced test set, while the remaining households were allocated for training. To emulate a realistic distributed scenario, we partitioned the training data across 10 fictitious energy utility companies, with no overlap. We followed the

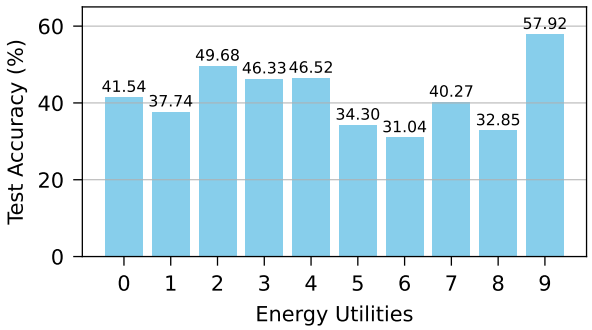


Fig. 3. Test accuracy over ten energy utilities in Siloed Training.

partitioning procedure of [38] to assign households to each utility based on class ratios sampled from a Dirichlet distribution with a concentration parameter of 0.3. This procedure generates statistically heterogeneous partitions, with each utility dataset holding a different number of households and an imbalanced class distribution. This setup replicates the characteristics of real-world distributed environments, where local energy utilities operate across diverse geographic regions and demographic segments.

B. Training Scenarios

For evaluation, we considered three training scenarios:

1) *Siloed Training*: is a baseline setup in which each energy utility independently trains a model on its own data. This baseline aims to illustrate the limitations of isolated learning and is the reference point for quantifying the benefits of cross-party collaboration.
2) *Centralized Training*: represents the ideal setting in which data from all utilities is aggregated into a single repository for standard training. Although this approach is not feasible in practice due to regulatory constraints preventing data centralization, it offers a benchmark for the maximum achievable performance.
3) *Privacy-Preserving FL*: uses the FL framework secured with HE presented in Sec. III-B, allowing secure collaborative training among energy utilities.

C. Preprocessing, Model & Training Hyperparameters

Prior to training, we applied a preprocessing pipeline to structure hourly consumption timeseries in a format suitable for temporal pattern extraction. The timeseries were first segmented into non-overlapping monthly windows, then reshaped into tensors with 4 channels, each collecting one week of data. This representation enables the model to learn both short-term and long-term consumption patterns while ensuring computational efficiency.

The model architecture is a sequential 1D-CNN composed of two convolutional layers followed by two fully-connected layers. The first convolutional layer uses 32 filters with a kernel size of 7, while the second uses 64 filters with a kernel size of 3. Each convolutional operation is followed by a ReLU activation function and a max pooling layer with kernel size and stride of 2. The resulting feature maps are then refined by the first fully connected layer, which returns a 512-dimensional representation processed by a ReLU activation and the final classification layer. The same model was used for all training scenarios to ensure a fair comparison.

Regarding the training hyperparameters, we used the same setup for both *Siloed Training* and *Centralized Training*. Specifically, the models were trained for 20,000 iterations with the SGD optimizer, using learning rate 0.01, Nesterov momentum 0.9, and batch size 64. With *Privacy-Preserving FL*, the training was conducted over 1,000 synchronization rounds, each consisting of 10 local training iterations per party, using the same optimizer configuration as in the previous scenarios. For HE, we employed the CKKS scheme implemented by the TenSEAL library, with cryptographic parameters set as follows: polynomial modulus degree to 8192, scale to 40 bits, and coefficient modulus to 200 bits.

TABLE I
HOUSEHOLD CHARACTERIZATION IDENTIFICATION ACCURACY UNDER THREE TRAINING SCENARIOS.

Training Scenario	Test Accuracy
Siloed Training	41.82%*
Centralized Training	81.09%
Privacy-Preserving FL	77.92%

* Average over 10 energy utilities.

D. Results

The bar plot in Fig. 3 reports the classification accuracy under the *Siloed Training* scenario measured on the test set for each energy utility (each denoted by a number ranging from 0 to 9). The analysis of the collected results reveals two main insights. First, we observe a significant variability in performance among the utilities, with accuracy ranging from 31.04% to 57.92%. This variability stems from the inherent statistical heterogeneity of the local training datasets, which differ in size, class coverage, and distribution. This observation highlights that the quality and class representation of local data substantially affect the model performance, making the outcome highly dependent on the available data subset of each utility. Second, the accuracy of all models is well below acceptable levels for practical usage. Even the best-performing utility reaches only a mere 57.92% accuracy, indicating limited capacity to generalize across the full range of sociodemographic profiles and electricity consumption patterns. These findings emphasize the limitations of *Siloed Training*, which fails to provide sufficient classification quality due to the lack of collaborative data sharing across parties.

A comparison across the three training strategies further emphasizes the limitations of *Siloed Training*. As shown in Table I, the average classification accuracy obtained in this setting is 41.82%, confirming the weak generalization ability of models trained on fragmented datasets. In contrast, the *Centralized Training* scenario achieves a substantially higher accuracy of 81.09%. However, this result represents an ideal upper bound, which is not achievable in real-world contexts due to strict data privacy regulations that prohibit the sharing

of raw consumption data across organizational boundaries. In this context, *Privacy-Preserving FL* offers a valuable alternative. It achieves 77.92% accuracy, just 3.17% points below the centralized baseline, while ensuring secure collaboration thanks to data locality and HE protection. This result demonstrates that HE-based FL provides a practical and regulation-compliant path to achieving high-quality sociodemographic classification while preserving data sovereignty.

V. Conclusion

Our work presented a training framework for privacy-preserving household characterization identification from electricity consumption data. Our contribution lies in the integration of two key components: (i) a synthetic data generation pipeline capable of modeling fine-grained sociodemographic signatures in household energy consumption data, and (ii) a federated learning architecture featuring homomorphic encryption to ensure secure collaborative training. Our experimental analysis revealed that siloed training is a weak option, leading to suboptimal and inconsistent performance across energy utilities. In contrast, the collected results validate the effectiveness of federated learning, achieving accuracy levels close to traditional centralized training. By demonstrating the value of FL for sociodemographic profiling under data regulation constraints, our study offers a strong foundation for advancing secure and collaborative intelligence in distributed energy systems.

References

[1] C. Beckel, L. Sadamori, T. Staake, and S. Santini, "Revealing household characteristics from smart meter data," *Energy*, 2014.

[2] Y. Kiguchi, M. Weeks, and R. Arakawa, "Predicting winners and losers under time-of-use tariffs using smart meter data," *Energy*, 2021.

[3] M. Castangia, R. Sappa, A. A. Girmay, C. Camarda, E. Macii, and E. Patti, "Detection of anomalies in household appliances from disaggregated load consumption," in *2021 International Conference on Smart Energy Systems and Technologies (SEST)*. IEEE, 2021, pp. 1–6.

[4] V. von Loessl, "Smart meter-related data privacy concerns and dynamic electricity tariffs: Evidence from a stated choice experiment," *Energy policy*, 2023.

[5] X. Cheng, C. Li, and X. Liu, "A review of federated learning in energy systems," in *I&CPS Asia)*, 2022.

[6] C. De Vizia, D. Salvatore Schiera, A. Macii, E. Patti, and L. Bottaccioli, "A simulation framework for urban electric mobility based on limited widespread data and spatial information," *IEEE Trans. Intell. Transp. Syst.*, 2024.

[7] C. De Vizia, A. Macii, E. Patti, and L. Bottaccioli, "A hierarchical and modular agent-oriented framework for power systems co-simulations," *Energy Informatics*, 2022.

[8] M. Pekey, Y. D. Çelebi, C. Anıl, and A. Levi, "Private information inference of households from electricity consumption data," in *BalkanCom*, 2021.

[9] A. Roshan and D. Ganga, "Intelligent categorization and interactive mechanism for smart demand side management of residential consumers," *Sustain. Energy Grids Netw.*, 2024.

[10] J. Lin, J. Ma, and J. G. Zhu, "Estimation of household characteristics with uncertainties from smart meter data," *Int. J. Electr. Power Energy Syst.*, 2022.

[11] H. Fang, J.-W. Xiao, and Y.-W. Wang, "Self-training convolutional autoencoder for consumer characteristics identification with imbalance datasets," *Eng. Appl. Artif. Intell.*, 2023.

[12] H. Wen, X. Liu, M. Yang, B. Lei, C. Xu, and Z. Chen, "A novel approach for identifying customer groups for personalized demand-side management services using household socio-demographic data," *Energy*, 2024.

[13] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *AISTATS*, 2017.

[14] R. Shrestha, M. Mohammadi, S. Sinaei, A. Salcines, D. Pampliega, R. Clemente, A. L. Sanz, E. Nowroozi, and A. Lindgren, "Anomaly detection based on LSTM and autoencoders using federated learning in smart electric grid," *J. Parallel Distributed Comput.*, 2024.

[15] M. Abdel-Basset, N. Moustafa, and H. Hawash, "Privacy-preserved generative network for trustworthy anomaly detection in smart grids: A federated semisupervised approach," *IEEE Trans. Ind. Informatics*, 2023.

[16] M. Wen, R. Xie, K. Lu, L. Wang, and K. Zhang, "Feddetect: A novel privacy-preserving federated learning framework for energy theft detection in smart grid," *IEEE Internet Things J.*, 2022.

[17] M. N. Fekri, K. Grolinger, and S. Mir, "Distributed load forecasting using smart meter data: Federated learning with recurrent neural networks," *Int. J. Electr. Power Energy Syst.*, 2022.

[18] M. M. Badr, M. M. E. A. Mahmoud, Y. Fang, M. J. Abdulaal, A. J. Aljohani, W. Alasmary, and M. I. Ibrahem, "Privacy-preserving and communication-efficient energy prediction scheme based on federated learning for smart grids," *IEEE Internet Things J.*, 2023.

[19] Y. Wang, M. Jia, N. Gao, L. V. Krannichfeldt, M. Sun, and G. Hug, "Federated clustering for electricity consumption pattern extraction," *IEEE Trans. Smart Grid*, 2022.

[20] W. Chen, S. Bu, X. Zhang, Y. Tao, Y. Zhang, and Z. Han, "Semi-supervised federated analytics for heterogeneous household characteristics identification," *IEEE Trans. Smart Grid*, 2024.

[21] J. Geiping, H. Bauermeister, H. Dröge, and M. Moeller, "Inverting gradients - how easy is it to break privacy in federated learning?" in *NeurIPS 2020*, 2020.

[22] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *EUROCRYPT*, 1999.

[23] J. H. Cheon, A. Kim, M. Kim, and Y. S. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *ASIACRYPT*, 2017.

[24] V. Peluso, E. Malan, A. Calimera, and E. Macii, "Private tensor freezing for an efficient federated learning with homomorphic encryption," in *ICCD*. IEEE, 2024.

[25] C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu, "Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning," in *ATC*, 2020.

[26] A. Benaissa, B. Retiat, B. Cebere, and A. E. Belfedhal, "Tenseal: A library for encrypted tensor operations using homomorphic encryption," in *DPML ICLR Workshop*, 2021.

[27] C. He, G. Liu, S. Guo, and Y. Yang, "Privacy-preserving and low-latency federated learning in edge computing," *IEEE Internet Things J.*, 2022.

[28] C. Hu and B. Li, "Maskcrypt: Federated learning with selective homomorphic encryption," *IEEE Trans. Dependable Secur. Comput.*, 2025.

[29] E. Malan, V. Peluso, A. Calimera, and E. Macii, "Communication-efficient federated learning with gradual layer freezing," *IEEE Embed. Syst. Lett.*, 2023.

[30] E. Malan, V. Peluso, A. Calimera, E. Macii, and P. Montuschi, "Automatic layer freezing for communication efficiency in cross-device federated learning," *IEEE Internet Things J.*, 2024.

[31] Eurostat, "Harmonised European Time Use Surveys (HETUS) 2018 Guidelines," Publications Office of the European Union,, Tech. Rep., 2020.

[32] J. Janssen and R. Manca, *Semi-Markov risk models for finance, insurance and reliability*. Springer Science & Business Media, 2007.

[33] J. D. Rhodes, W. J. Cole, C. R. Upshaw, T. F. Edgar, and M. E. Webber, "Clustering analysis of residential electricity demand profiles," *Applied Energy*, 2014.

[34] ISTAT, "Multiscopo sulle famiglie: uso del tempo," 2013. [Online]. Available: https://www.istat.it/microdati/multiscopo-sulle-famiglie-uso-del-tempo/

[35] ——, "Annuario statistico italiano," ISTAT, Tech. Rep., 2022.

[36] ARENA, "Demand response customer insights report," 2018. [Online]. Available: https://arena.gov.au/assets/2018/08/demand-response-consumer-insights-report.pdf

[37] T. Yunusov and J. Torriti, "Distributional effects of time of use tariffs based on electricity demand and time use," *Energy Policy*, 2021.

[38] H. Wang, M. Yurochkin, Y. Sun, D. S. Papailiopoulos, and Y. Khazaeni, "Federated learning with matched averaging," in *ICLR*, 2020.