

FedCCL: Federated Clustered Continual Learning Framework for Privacy-focused Energy Forecasting

Michael A. Helcig
Distributed Systems Group
TU Wien
michael.helcig@student.tuwien.ac.at

Stefan Nastic
Distributed Systems Group
TU Wien
snastic@dsg.tuwien.ac.at

Abstract—Privacy-preserving distributed model training is crucial for modern machine learning applications, yet existing Federated Learning approaches struggle with heterogeneous data distributions and varying computational capabilities. Traditional solutions either treat all participants uniformly or require costly dynamic clustering during training, leading to reduced efficiency and delayed model specialization. We present FedCCL (Federated Clustered Continual Learning), a framework specifically designed for environments with static organizational characteristics but dynamic client availability. By combining static pre-training clustering with an adapted asynchronous FedAvg algorithm, FedCCL enables new clients to immediately profit from specialized models without prior exposure to their data distribution, while maintaining reduced coordination overhead and resilience to client disconnections. Our approach implements an asynchronous Federated Learning protocol with a three-tier model topology — global, cluster-specific, and local models — that efficiently manages knowledge sharing across heterogeneous participants. Evaluation using photovoltaic installations across central Europe demonstrates that FedCCL’s location-based clustering achieves an energy prediction error of 3.93% ($\pm 0.21\%$), while maintaining data privacy and showing that the framework maintains stability for population-independent deployments, with 0.14 percentage point degradation in performance for new installations. The results demonstrate that FedCCL offers an effective framework for privacy-preserving distributed learning, maintaining high accuracy and adaptability even with dynamic participant populations.

Index Terms—Federated Learning, Clustered Federated Learning, Energy Forecasting, Privacy-Preserving Computing, Distributed Machine Learning, Time Series Analysis, Smart Grid, Renewable Energy, Edge Computing

I. INTRODUCTION

The Federated Learning (FL) paradigm [1], [2] has emerged as a pivotal solution for privacy-preserving machine learning, enabling multiple participants to collaboratively train models while maintaining data privacy. This approach has found significant applications across domains where sensitive information must remain decentralized, such as healthcare, finance, and industrial systems. Traditional FL frameworks operate through synchronous rounds of communication, where participants collectively train a single shared model. However, this approach often proves inefficient and inflexible, particularly when dealing with heterogeneous data distributions and varying computational capabilities across participants [3], [4].

Current challenges in FL deployments include balancing model specialization, meaning the ability to excel at specific

client scenarios, with overall generalization, which ensures robust performance across diverse client populations. While many environments have stable organizational characteristics (such as geographic location or hardware specifications), they face dynamic challenges in client availability and connectivity, another issue currently explored [5], [6]. Existing solutions typically address static or dynamic aspects independently, but not both.

In terms of communication, synchronous FL approaches [7] require all participants to complete their local training before model aggregation, leading to potential bottlenecks and reduced system efficiency when participants operate at different speeds or face connectivity issues. Asynchronous FL algorithms [3], [8], [9] address these limitations by enabling independent model updates. The challenge of non-iid (non-independent and identically distributed) data has sparked significant research in Clustered Federated Learning (CFL) approaches. Recent frameworks like Iterative Clustered FL (ICFL) [10] propose dynamic clustering mechanisms based on model behavior during training. Weight-driven approaches [11] exploit correlations between local model weights and data distributions to form clusters. However, these methods typically require multiple training rounds to establish stable clusters and may not immediately benefit new participants joining the federation.

Previous research has extensively explored both solar prediction methodologies [12]–[14] as well as the application of FL. Approaches range from hierarchical clustering for residential load forecasting [15] to fuzzy clustering for solar power generation [16]. Recent work has explored deep learning models for PV power forecasting [14], incorporating domain knowledge [17] and simplified LSTM [18] architectures [19]. There is also some work concerning pre-training clustering [20], but without a focus on transfer capabilities, asynchronous rounds and production forecasting. Load forecasting was addressed several times. [20]–[22]

This paper presents FedCCL (Federated Clustered Continual Learning), a framework that addresses these challenges through a combination of pre-training clustering and asynchronous Federated Learning. Unlike most of the existing approaches that perform clustering during or after training [23], [24], FedCCL employs DBSCAN clustering based on static characteristics before training begins. This approach enables

immediate model specialization while reducing coordination overhead. Furthermore, participants can belong to multiple clusters simultaneously, facilitating more nuanced knowledge sharing than strict partitioning approaches [16]. This paper’s main contributions include:

- **FedCCL Framework:** A Federated Learning framework that integrates clustered pre-training with an enhanced asynchronous FedAvg algorithm. The framework operates through a two-phase approach, initially clustering clients based on their inherent system properties before training, followed by client-driven updates with model locking during aggregation. Mitigating the performance degradation typically seen in asynchronous Federated Learning with heterogeneous data while maintaining reduced overhead.
- **FedCCL Predict & Evolve:** Through our system property-based clustering approach, FedCCL creates a framework that provides a specialized model for newly joining clients without requiring prior exposure to their specific data distributions. In the "Predict" phase, new clients can immediately benefit from these highly specialized models to generate predictions. As clients begin contributing their own data, they enter the "Evolve" phase, where they participate in training and refining cluster-specific models. Our evaluation demonstrates this capability through robust generalization metrics, where models achieve nearly identical performance levels for both training and independent populations, with mean error rates showing minimal degradation of only 0.14 percentage points for new installations.
- **FedCCL Case Study:** We validate our framework through implementation in solar production forecasting, using real-world data from sites across central Europe, clustered by geographical location and panel orientation. Our experimental results demonstrate superior accuracy compared to both traditional asynchronous Federated Learning and centralized approaches, confirming its effectiveness. In our solar production forecasting implementation using real-world data, the framework with location-based clustering achieved mean errors of 6.44% ($\pm 0.17\%$) for power prediction and 3.93% ($\pm 0.21\%$) for energy prediction, demonstrating superior accuracy compared to the centralized baseline which had higher errors of 6.73% ($\pm 0.14\%$) and 4.18% ($\pm 0.18\%$) respectively.

This paper is organized into six sections. We begin by introducing the problem and motivation (Section I), followed by a detailed description of the FedCCL framework and its clustered continual learning approach (Section II). We then present an energy forecasting case study (Section III) with experimental results that validate our methodology (Section IV). The paper concludes with a review of related work (Section V) and a discussion of future research directions. (Section VI)

II. FEDCCL FRAMEWORK

The proposed FedCCL framework is designed as a centralized and horizontal system [1], it uses a central server to coordinate learning while focusing on scenarios where clients share the same feature space. This section outlines the framework’s architecture and operational workflow.

A. Architecture Overview

FedCCL operates on a three-tier model hierarchy (as illustrated in Figure 1) that balances global knowledge sharing with specialized learning. At the highest level, a global model (sparsely-dashed) synthesizes knowledge across all clients, capturing common patterns and general features applicable to the entire system. The middle tier consists of cluster-specific models (dashed), each specialized for a group of similar clients. These cluster models, labeled as Cluster Model A through N, refine the global knowledge to better serve the particular characteristics and requirements of their respective client groups, enabling more targeted and efficient learning. At the lowest tier, local models on individual clients (dotted) provide fine-grained specialization, allowing each participant (Client 1 through Client n) to adapt to their specific needs while maintaining the privacy of their local data. This hierarchical structure ensures that the global model captures broad trends, cluster models address group-specific patterns, and local models fine-tune these insights to excel in client-specific tasks, thereby mitigating challenges posed by non-iid data.




Fig. 1. Model Hierarchy of FedCCL

B. Pre-training Clustering

One of FedCCL’s distinctive features is its pre-training clustering phase, which precedes model training. Unlike traditional FL approaches, FedCCL employs DBSCAN [25] to establish a stable cluster structure based on invariant client characteristics such as geographic location, hardware specifications, or operational parameters. This density-based approach naturally identifies optimal client groupings while filtering outliers. The framework can be extended with the incremental DBSCAN variant [26], enabling network expansion without disrupting established structures. This Predict & Evolve capability of FedCCL allows new clients to immediately benefit from specialized models tailored to their cluster’s characteristics

without requiring complete reclustering. By establishing these relationships before training, FedCCL creates a foundation for targeted model specialization while minimizing computational overhead typically associated with dynamic clustering approaches.

C. Asynchronous Training Process

The core of FedCCL implements a novel asynchronous FL approach, as detailed in Algorithm 1. After initialization (lines 1-2), each client operates independently and in parallel (line 3). During training, clients first update their local models using their private datasets (lines 5-6). Following the local update, they participate in parallel cluster-specific training sessions (lines 7-11), where they request the current cluster model state, perform local training, compute update deltas, and submit their contributions. Finally, clients engage with the global model in a similar manner (lines 12-15). The server handles these asynchronous updates through a dedicated handler function (lines 19-25) that locks the specific model during aggregation to ensure consistency while maximizing parallel operations. In contrast to synchronous FL, this design removes barriers, making it particularly suitable for settings with varying client network conditions.

Algorithm 1: FedCCL Training Process

Data: K clients with local datasets D_k
Result: Trained models across hierarchy levels

- 1 Initialize server model store with global and cluster models;
- 2 Initialize local model an set of cluster keys $client_clusters$ for each client;
- 3 **for** each *distributed client* k **in parallel** **do**
- 4 **while** training **do**
- 5 Preprocess local data D_k ;
- 6 // Train local model
 $w_{k_{local}} \leftarrow \text{TrainModel}(w_{local}, D_k)$;
- 7 // Train cluster models
for $cluster_key$ **in** $client_clusters$ **in parallel** **do**
- 8 $w \leftarrow \text{RequestModel}(level = \text{cluster}, cluster_key)$;
- 9 $w_k \leftarrow \text{TrainModel}(w, D_k)$;
- 10 $\delta_k \leftarrow \text{ComputeModelMetaDelta}()$;
- 11 $\text{HandleModelUpdate}(level = \text{cluster}, cluster_key = cluster_key, w_{updated} = w_k, \delta_{new} = \delta_k)$;
- 12 **end**
- 13 // Train global model
 $w \leftarrow \text{RequestModel}(level = \text{global})$;
- 14 $w_k \leftarrow \text{TrainModel}(w, D_k)$;
- 15 $\delta_k \leftarrow \text{ComputeModelDelta}()$;
- 16 $\text{HandleModelUpdate}(level = \text{'global'}, w_{updated} = w_k, \delta_{new} = \delta_k)$;
- 17 **end**
- 18 **end**
- 19 // Server Update Handler
Function $\text{HandleModelUpdate}(level, cluster_key, w_{updated}, \delta_{new})$:
 $m \leftarrow \text{GetModel}(level, cluster_key)$;
- 21 **if** $\text{AcquireLock}(m)$ **then**
- 22 $m \leftarrow \text{AggregateModels}(m, w_{updated}, \delta_{new})$;
- 23 $\text{UpdateModelStore}(level, cluster_key, m)$;
- 24 $\text{ReleaseLock}(m)$;
- 25 **end**

D. Model Aggregation Strategy

FedCCL builds upon FedAvg’s [7] weighted averaging approach detailed in Algorithm 2. The aggregation process, performed entirely on the server, first evaluates whether the model has been updated by another client by checking if the difference in rounds is exactly 1 (line 1). If this condition is met, it returns the updated model without aggregation (line 2). When updates are non-sequential, the algorithm computes the total number of training samples (line 4) and performs a layer-wise weighted aggregation (lines 7-11). Each client’s influence is proportional to the number of new samples learned in that training round, relative to the total samples processed, determined through ratios calculated for both the base model and new updates. The framework then updates essential metadata (lines 11-13), including cumulative sample count, training epochs, and federation rounds, before returning the aggregated model (line 14). This metadata enables systematic tracking of the federation’s learning progress and ensures proper model versioning.

Algorithm 2: FedCCL Model Aggregation

Function $\text{AggregateModels}(w_{base}, w_{updated}, \delta_{new})$:
// Check if updates are sequential (i.e. round of updated model is exactly one ahead of current)

- 1 **if** $w_{updated}.round == w_{base}.round + 1$ **then**
- 2 **return** $w_{updated}$;
- 3 **end**
- 4 $samples_total \leftarrow w_{base}.samples_learned + w_{updated}.samples_learned$;
- 5 Initialize $w_{aggregated}$ with zeros;
- 6 **for** each weight layer i **do**
- 7 $ratio_{base} \leftarrow \frac{w_{base}.samples_learned}{samples_total}$;
- 8 $ratio_{new} \leftarrow \frac{w_{updated}.samples_learned}{samples_total}$;
- 9 $w_{aggregated}[i] \leftarrow w_{base}[i] \times ratio_{base} + w_{updated}[i] \times ratio_{new}$;
- 10 **end**
- 11 // Update model metadata
 $meta.samples_learned \leftarrow w_{base}.samples_learned + \delta_{new}.samples_learned$;
- 12 $meta.epochs_learned \leftarrow w_{base}.epochs_learned + \delta_{new}.epochs_learned$;
- 13 $meta.round \leftarrow w_{base}.round + \delta_{new}.round$;
- 14 **return** $\text{ModelData}(meta, w_{aggregated})$;

E. Continual Learning and Convergence

FedCCL implements an asynchronous and continuous learning approach where clients can join or leave the training process at any time, and the system can theoretically continue learning indefinitely. Unlike traditional Federated Learning systems that operate in synchronized rounds, FedCCL’s asynchronous nature allows clients to independently retrieve the current model state, perform local training, and contribute updates back to the system at their own pace. This flexibility eliminates the need for coordination between clients and enables continuous model improvement without predetermined stopping points. FedCCL addresses the challenge of

catastrophic forgetting using a regularization-based approach [27]. This method, often referred to as L2 regularization, penalizes deviations from important parameters of previously learned tasks. By constraining significant weights to remain close to their prior values, the framework ensures that newly learned information does not overwrite essential knowledge. While round-based systems provide more predictable convergence patterns, FedCCL’s asynchronous updates introduce non-determinism in the training process.

F. Privacy Preservation

FedCCL ensures privacy by keeping raw data on client devices, sharing only model parameters and essential metadata through secure channels. The clustering mechanism adds privacy by grouping similar clients, making individual inference more difficult. These features enable deployment in sensitive domains like healthcare, finance, and industrial systems.

G. Implementation

The framework was implemented following production engineering practices with containerization, emphasizing scalability and security for real-world edge computing scenarios. All source code, datasets, and analysis scripts are publicly available¹.

III. CASE STUDY: ENERGY FORECASTING WITH FEDCCL

Our case study implements and evaluates FedCCL using a real-world scenario for solar production forecasting of individual households and commercial & industrial (C&I) installations across central Europe. Such predictions are critical in the transition towards renewable energy systems, especially for use-cases such as battery charging, heating, and electric vehicle charging during periods of peak energy availability.

The clustering phase in our study follows two primary objectives. First, we implement location-based clustering, as illustrated in Figure 2, where installations are grouped based on their geographical proximity.




Fig. 2. Location-based Clustering (Coordinates Hidden)

This approach exploits the strong correlation between regional weather patterns and solar production. The figure demonstrates the distribution of three main clusters (marked in blue, orange, and green), with test sites (marked with red crosses) selected as the installations closest to each cluster’s

¹<https://github.com/polaris-slo-cloud/fedccl>

centroid. For privacy preservation, the exact longitude and latitude coordinates are intentionally obscured. Second, we incorporate installation-specific characteristics by clustering based on panel orientation, determined through analysis of peak power generation patterns during clear-sky conditions. This dual-clustering strategy ensures the framework can effectively account for both regional weather patterns and site-specific performance characteristics.

A. Dataset and Forecasting Objective

The dataset spans approximately fifteen months, comprising solar production data recorded at 15-minute intervals (96 measurements per day) and complementary hourly weather forecast data, which includes essential meteorological parameters affecting solar production. To harmonize these different temporal resolutions, weather forecasts are duplicated across their respective 15-minute intervals.

Figure 3 illustrates a representative week of normalized data, demonstrating the relationship between meteorological conditions and power generation. The black line represents the actual power output. The daily production cycles are clearly visible, with peak generation typically occurring during midday hours when solar radiation is strongest. The influence of weather conditions is evident in the varying production patterns: days with high cloud coverage (purple) show reduced and more irregular power generation, while clear days exhibit smooth, bell-shaped production curves. Snow depth (green) and precipitation (red) remain at zero during this period. The training window for predictions encompasses seven days of historical data. Based on this historical data and weather forecasts, the model predicts power generation for the next 24 hours in 15-minute intervals, providing 96 discrete forecasting points per day.




Fig. 3. Data Shape

B. Feature Selection and Processing

We identified key features that impact prediction accuracy. The most influential weather-related parameters include ground-level solar radiation and Global Horizontal Irradiance (GHI), which directly correlate with potential energy production. Snow depth and precipitation levels are incorporated to account for weather conditions that may obstruct panel efficiency, while cloud coverage data helps identify periods of potentially irregular production. Temporal features are used for capturing both daily and seasonal patterns in solar energy

production. The system tracks the minute of day and day of year, both normalized to capture cyclical patterns in solar availability. All features undergo normalization to ensure consistent training across sites with varying capacities. Weather parameters are scaled using regional maximum values typical for central Europe, while production values are normalized relative to each site’s maximum capacity (kWp).

TABLE I
FEATURE DESCRIPTIONS

Feature Name	Description	Range (min. - max.)
<i>solar_rad</i>	Ground-level solar radiation (W/m ²)	0 - 956.2 W/m ²
<i>ghi</i>	Global Horizontal Irradiance measured outside atmosphere (W/m ²)	0 - 956.21 W/m ²
<i>snow_depth</i>	Snow depth affecting panel efficiency (mm)	0 - 1178.6 mm
<i>precip</i>	Precipitation levels (mm)	0 - 14.78 mm
<i>clouds</i>	Cloud coverage (high-coverage days filtered)	0 - 100.0 %
<i>minute_of_day</i>	Temporal feature capturing daily cycles (minutes)	0 - 1440 minutes
<i>day_of_year</i>	Temporal feature capturing seasonal cycles (days)	0 - 365 days

IV. EVALUATION

A. Evaluation Objective and Baselines

We evaluate FedCCL against centralized baselines (CentralizedAll with complete data access and CentralizedContinual with progressive data availability) to demonstrate that we achieve similar or better performance and that clustering provides benefits over a globally shared model. Each client trains four models: local, federated global, and two federated cluster models (location-based and orientation-based). Using an 80-20 training-testing split across 100 experimental runs, we measure performance during both 24-hour periods and daytime operations (06:00-21:00) to ensure statistical significance and context-appropriate evaluation in real-world scenarios.

B. Performance Metrics

We assess prediction accuracy using two distinct Mean Absolute Percentage Error (MAPE) metrics. For instantaneous power predictions, we normalize errors against each system’s rated capacity (kilowatts peak, kWp):

$$\text{Power Error} = \left(\frac{|\text{predicted} - \text{actual}|}{\text{kWp}} \right) \times 100\%$$

For cumulative energy we normalize errors against the theoretical daily maximum production, calculated as 12 hours at rated capacity:

$$\text{Energy Error} = \left(\frac{|\text{predicted} - \text{actual}|}{\text{kWp} \times 12} \right) \times 100\%$$

This normalization enables fair comparison across solar installations regardless of their size. We evaluate performance using mean and maximum values of *error power* for instantaneous predictions and *error energy* for cumulative predictions, allowing us to assess both typical performance and edge cases. Detailed results are presented in Table II.

C. Baseline Performance Analysis

The CentralizedAll model, with complete initial data access, establishes our primary baseline with a mean error for power prediction of 6.24% ($\pm 0.07\%$). This relatively low standard deviation indicates highly consistent performance across different runs. During daytime hours (6:00-21:00), the mean error increases to 9.27%. The model achieves a mean error for energy prediction of 3.46%, demonstrating better accuracy in cumulative energy prediction compared to instantaneous power forecasting.

The CentralizedContinual model, which mirrors real-world data availability constraints, shows degraded performance with a mean error power of 6.73% ($\pm 0.14\%$). This 0.49 percentage point increase in error compared to the CentralizedAll model quantifies the impact of progressive data availability on prediction accuracy. The higher standard deviation also indicates less consistent performance across runs.

D. Federated Model Performance Analysis

1) *Global Model*: The Federated Global model achieves a mean error power of 6.59% ($\pm 0.24\%$), positioning its performance between the two centralized baselines. The increased standard deviation suggests more variability in the accuracy of the prediction, likely due to the asynchronous nature of federated updates. Notably, the model’s maximum error power of 70.34% represents a significant improvement over the CentralizedContinual model’s 84.31%, indicating better handling of extreme cases.

2) *Location-Based Clustering*: The Location-based clustering approach demonstrates superior performance among federated models, achieving a mean error power of 6.44% ($\pm 0.17\%$). This represents only a 0.20 percentage point degradation compared to the CentralizedAll baseline while maintaining complete data privacy. The model’s daytime performance (9.56% mean error) shows particular strength in handling peak production periods, as illustrated by the exemplary prediction shown in Figure 4.

3) *Orientation-Based Clustering*: The Orientation-based clustering model shows slightly diminished performance with a mean error power of 6.54% ($\pm 0.09\%$). While this remains competitive, the smaller improvement over the Global model suggests that panel orientation alone may not be as strong a differentiator for model specialization as geographical location. However, this model demonstrates the highest consistency across runs, as indicated by its low standard deviation.

E. Population Independent Performance Analysis

A key aspect of our evaluation, namely the FedCCL Predict & Evolve contribution, focuses on the framework’s ability to serve new installations without prior training data. This analysis reveals several key insights:

The Global and Location models demonstrate robust generalization capabilities, achieving mean error power rates of 6.60% and 6.58% respectively for independent sites. These results represent minimal degradation from their performance on the training population (0.01 and 0.14 percentage points),

TABLE II
COMPREHENSIVE MODEL PERFORMANCE COMPARISON (ALL VALUES IN PERCENTAGES, BASED ON 100 RUNS)

Metric	Centralized (All)	Centralized (Continual)	Federated Global	Federated Location	Federated Orientation	Federated Local
Overall Performance						
Mean Error Power	6.24 ± 0.07	6.73 ± 0.14	6.59 ± 0.24	6.44 ± 0.17	6.54 ± 0.09	6.34 ± 0.07
Max Error Power	61.60 ± 1.56	84.31 ± 1.10	70.34 ± 1.70	71.00 ± 1.70	70.80 ± 1.96	83.90 ± 1.06
Mean Error Energy	3.46 ± 0.08	4.18 ± 0.18	4.17 ± 0.44	3.93 ± 0.21	4.05 ± 0.19	4.19 ± 0.12
Daytime Performance (6:00-21:00)						
Mean Error Day Power	9.27 ± 0.11	10.01 ± 0.21	9.75 ± 0.36	9.56 ± 0.26	9.71 ± 0.13	9.41 ± 0.11
Mean Error Day Energy	4.38 ± 0.10	5.26 ± 0.21	5.25 ± 0.52	4.96 ± 0.25	5.16 ± 0.23	5.33 ± 0.15




Fig. 4. Location-based Cluster Model Performance Example

indicating strong transferability of learned patterns. The Orientation model shows less effective generalization with a mean error power of 9.13% for independent sites, a 2.59 percentage point increase from its training population performance. Illustrated by the exemplary prediction shown in Figure 5.

During daytime hours, the performance difference becomes more pronounced, with the Location model maintaining a mean error day power of 9.72% compared to the Orientation model's 13.73%. This divergence emphasizes the superior generalization capabilities of location-based clustering for new installations.

F. Energy vs. Power Prediction Analysis

Energy prediction consistently demonstrates lower error rates than power prediction across all models. The CentralizedAll model achieves 3.46% mean energy error, while federated approaches range from 3.93% to 4.19%. Since energy is calculated through integration of power over time, this improved accuracy stems from an averaging effect that diminishes instantaneous prediction errors. This pattern persists in




Fig. 5. Population Independent Performance Example

independent populations, with the Location model achieving 4.13% mean energy error. This particularly benefits day-ahead planning and resource allocation applications.

G. Implications for Real-World Deployment

FedCCL provides a real-world applicable system delivering centralized-level performance while preserving privacy. The Location model (6.44% error) outperforms the Global model (6.59%), with minimal degradation (< 0.15 percentage points) in new deployments, offering immediate value across various application domains without compromising privacy.

V. RELATED WORK

Centralized solar production forecasting has been extensively studied in the literature, with comprehensive reviews [12]–[14] documenting significant advances in prediction methodologies. Long Short-Term Memory (LSTM) networks [18], which we also employ in our FedCCL Case Study, have established themselves as a leading approach for time series prediction in this domain. Their effectiveness has been validated across multiple studies [28]–[30], consistently demon-

strating robust performance in handling temporal dependencies inherent in solar production. Suresh et al. [31] achieved 2.6% MAE relative to peak capacity for hour-ahead forecasting. While their shorter prediction horizon differs from our multi-day weather-based forecasts, our evaluation shows comparable error magnitudes. While centralized approaches have shown promising results, they face fundamental limitations in scenarios requiring data privacy. The emergence of privacy-preserving distributed learning in the energy domain has created three distinct categories of solutions, each addressing different aspects: traditional FL approaches, clustering-enhanced methods, and domain-specific optimizations.

Traditional Federated Learning Approaches. Traditional FL implementations in energy forecasting, such as work from Stanford University [32] and STANN for GHI forecasting [33], as well recent work exploring the other side, i.e., load forecasting. [21], [22] demonstrate the fundamental viability of FL in that domain. However, these approaches rely on synchronous training rounds and single global models, which create two significant limitations. First, the synchronous nature requires all participants to complete their updates before proceeding, creating potential bottlenecks when participants operate at different speeds or face connectivity issues. Second, the single global model struggles to capture local variations effectively. FedCCL addresses these limitations through its asynchronous architecture and multi-tiered model hierarchy.

Clustering-Enhanced Methods. More sophisticated approaches have emerged that incorporate clustering to improve prediction accuracy. The FL+HC system for residential load forecasting [15] demonstrate that grouped learning can improve model performance by 5% or more. However, this method establish cluster after N rounds of training and clients belong to one cluster only. Similarly, fuzzy clustered FL for solar generation [16] also relies on cluster generation during training. Some approaches, like Tun et al.’s predefined clustering system [20], move closer to FedCCL’s philosophy but still rely on synchronous updates. FedCCL differs by performing clustering before training based on static problem specific characteristics and using asynchronous communication, better suited for real-world scenarios.

Domain-Specific Optimizations. Recent work has focused on incorporating domain knowledge into FL architectures, particularly for solar forecasting. Physics-constrained approaches [17] and specialized deep learning architectures [14] demonstrate the importance of domain expertise. Edge computing implementations [34] have shown promise in reducing communication overhead but typically maintain traditional FL patterns like a single global model and synchronous rounds. FedCCL builds upon these insights by combining domain-aware clustering (such as geographic location and panel orientation) with flexible client participation.

While recent approaches have explored individual elements - such as clustering for FL [15], [16], [20], domain-specific optimizations [17], [34], and basic privacy-preserving frameworks [32], [33] - no existing solution combines these aspects effectively. Current clustering methods require dynamic

updates during training, domain-specific approaches remain limited by synchronous operations, and traditional implementations cannot capture local variations. FedCC uniquely integrates these components through pre-training clustering and asynchronous updates, achieving accuracy comparable to centralized approaches [31] while maintaining privacy and enabling immediate model specialization.

VI. CONCLUSION & FUTURE WORK

This paper presents FedCCL, a framework that combines pre-training clustering with asynchronous FL. Through comprehensive evaluation in solar energy forecasting, FedCCL demonstrates performance that matches and exceeds centralized approaches while preserving data privacy and enabling immediate value delivery through pre-training clustering. The three main contributions — the FedCCL Framework, FedCCL Predict & Evolve, and the FedCCL Case Study — establish both theoretical foundations and practical viability. The location-based clustering approach achieved 6.44% mean error in power prediction, outperforming the CentralizedContinual baseline by 0.29 percentage points while maintaining data privacy. The framework’s robust performance in population-independent scenarios further validates its usefulness for deployment to new installations. We intend to extend the FedCCL framework in multiple directions by investigating cluster granularity and intersection handling, specifically examining the impact of hierarchical sub-clusters on model performance and developing optimal strategies for managing overlapping cluster characteristics. Furthermore, defining definite criteria which model to use in the inference phase. Lastly, areas for exploration encompass framework generalization to other domains with inherent clustering properties and analysis of convergence characteristics. Given the framework’s demonstrated effectiveness in solar forecasting, these research directions promise to advance the foundations of privacy-preserving distributed learning while enabling new industrial applications.

VII. ACKNOWLEDGEMENTS

This work is partially funded by the Austrian Research Promotion Agency (FFG) as part of the RapidREC project (No. 903884). This research was funded by the EU’s Horizon Europe Research and Innovation Program as part of the NexaSphere project (GA No. 101192912). This work has received funding from the Austrian Internet Stiftung under the NetIdee project LEOTrek (ID 7442). The data has been generously provided by neoom AG.

REFERENCES

- [1] Q. Yang, Y. Liu, T. Chen, and Y. Tong, “Federated Machine Learning: Concept and Applications,” *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 1–19, Mar. 2019. [Online]. Available: <https://dl.acm.org/doi/10.1145/3298981>
- [2] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, “Federated Optimization: Distributed Machine Learning for On-Device Intelligence,” 2016, version Number: 1. [Online]. Available: <https://arxiv.org/abs/1610.02527>

- [3] C. Xu, Y. Qu, Y. Xiang, and L. Gao, "Asynchronous federated learning on heterogeneous devices: A survey," *Computer Science Review*, vol. 50, p. 100595, Nov. 2023. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S157401372300062X>
- [4] R. Gajjanin, A. Danilenka, A. Morichetta, and S. Nastic, "Towards adaptive asynchronous federated learning for human activity recognition," in *Proceedings of the 14th International Conference on the Internet of Things (IoT 2024)*. New York, NY, USA: Association for Computing Machinery, 2024.
- [5] P. M. Mammen, "Federated Learning: Opportunities and Challenges," 2021, version Number: 1. [Online]. Available: <https://arxiv.org/abs/2101.05428>
- [6] S. Bharati, M. R. H. Mondal, P. Podder, and V. S. Prasath, "Federated learning: Applications, challenges and future directions," *HIS*, vol. 18, no. 1-2, pp. 19–35, May 2022. [Online]. Available: <https://journals.sagepub.com/doi/full/10.3233/HIS-220006>
- [7] H. Brendan McMahan, H. B. McMahan, Eider B Moore, E. Moore, Daniel Ramage, D. Ramage, Seth Hampson, S. Hampson, Blaise Agüera y Arcas, and B. Agüera y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," *arXiv: Learning*, Feb. 2016, arXIV_ID: 1602.05629 MAG ID: 2950745363 S2ID: d1dbf643447405984eeef098b1b320dee0b3b8a7.
- [8] Cong Xie, C. Xie, Oluwasanmi Koyejo, O. Koyejo, S. Koyejo, Indranil Gupta, and I. Gupta, "Asynchronous federated optimization," *arXiv: Distributed, Parallel, and Cluster Computing*, Mar. 2019, arXIV_ID: 1903.03934 MAG ID: 2921434559 S2ID: 8f6602b6be2962dacf0b563e73852183e628ddf.
- [9] Y. Chen, Y. Ning, M. Slawski, and H. Rangwala, "Asynchronous Online Federated Learning for Edge Devices with Non-IID Data," in *2020 IEEE International Conference on Big Data (Big Data)*. Atlanta, GA, USA: IEEE, Dec. 2020, pp. 15–24. [Online]. Available: <https://ieeexplore.ieee.org/document/9378161/>
- [10] Y. Yan, X. Tong, and S. Wang, "Clustered Federated Learning in Heterogeneous Environment," *IEEE Trans. Neural Netw. Learning Syst.*, vol. 35, no. 9, pp. 12796–12809, Sep. 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10106044/>
- [11] Md Sirajul Islam, Simin Javaherian, Fei Xu, Xu Yuan, Li Chen, and Nianfeng Tzeng, "FedClust: Optimizing Federated Learning on Non-IID Data Through Weight-Driven Client Clustering," *IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum*, 2024, arXIV_ID: 2403.04144 S2ID: e83384370963e3901c5d268a8ba6383527006868.
- [12] P. Kumari and D. Toshniwal, "Deep learning models for solar irradiance forecasting: A comprehensive review," *Journal of Cleaner Production*, vol. 318, p. 128566, Oct. 2021. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0959652621027736>
- [13] X. Qing and Y. Niu, "Hourly day-ahead solar irradiance prediction using weather forecasts by LSTM," *Energy*, vol. 148, pp. 461–468, Apr. 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0360544218302056>
- [14] J. Yu, X. Li, L. Yang, L. Li, Z. Huang, K. Shen, X. Yang, X. Yang, Z. Xu, D. Zhang, and S. Du, "Deep Learning Models for PV Power Forecasting: Review," *Energies*, vol. 17, no. 16, p. 3973, Aug. 2024. [Online]. Available: <https://www.mdpi.com/1996-1073/17/16/3973>
- [15] Christopher Briggs, C. C. Briggs, Zhong Fan, Zhong Fan, Peter Andras, and Péter András, "Federated Learning for Short-term Residential Load Forecasting," *IEEE open access journal of power and energy*, pp. 1–1, Jan. 2022, arXIV_ID: 2105.13325 MAG ID: 4295308572 S2ID: 3bdf502bcd49f4077dcd53a75f35cac7b0fc0b2c.
- [16] E. Yoo, H. Ko, and S. Pack, "Fuzzy Clustered Federated Learning Algorithm for Solar Power Generation Forecasting," *IEEE Trans. Emerg. Topics Comput.*, vol. 10, no. 4, pp. 2092–2098, Oct. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9687124/>
- [17] X. Luo, D. Zhang, and X. Zhu, "Deep learning based forecasting of photovoltaic power generation by incorporating domain knowledge," *Energy*, vol. 225, p. 120240, Jun. 2021. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0360544221004898>
- [18] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997. [Online]. Available: <https://direct.mit.edu/neco/article/9/8/1735-1780/6109>
- [19] C.-H. Liu, J.-C. Gu, and M.-T. Yang, "A Simplified LSTM Neural Networks for One Day-Ahead Solar Power Forecasting," *IEEE Access*, vol. 9, pp. 17174–17195, 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9333638/>
- [20] Y. L. Tun, K. Thar, C. M. Thwal, and C. S. Hong, "Federated Learning based Energy Demand Prediction with Clustered Aggregation," in *2021 IEEE International Conference on Big Data and Smart Computing (BigComp)*. Jeju Island, Korea (South): IEEE, Jan. 2021, pp. 164–167. [Online]. Available: <https://ieeexplore.ieee.org/document/9373194/>
- [21] Y. M. Saputra, D. T. Hoang, D. N. Nguyen, E. Dutkiewicz, M. D. Mueck, and S. Srikanthswara, "Energy Demand Prediction with Federated Learning for Electric Vehicle Networks," 2019, version Number: 1. [Online]. Available: <https://arxiv.org/abs/1909.00907>
- [22] L. Richter, S. Lenk, and P. Bretschneider, "Advancing Electric Load Forecasting: Leveraging Federated Learning for Distributed, Non-Stationary, and Discontinuous Time Series," *Smart Cities*, vol. 7, no. 4, pp. 2065–2093, Jul. 2024. [Online]. Available: <https://www.mdpi.com/2624-6511/7/4/82>
- [23] Felix Sattler, F. Sattler, Klaus-Robert Müller, K.-R. Müller, Wojciech Samek, and W. Samek, "Clustered Federated Learning: Model-Agnostic Distributed Multitask Optimization Under Privacy Constraints," *IEEE Transactions on Neural Networks*, vol. 32, no. 8, pp. 3710–3722, 2020, mag ID: 3080934299.
- [24] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, "An Efficient Framework for Clustered Federated Learning," *IEEE Trans. Inform. Theory*, vol. 68, no. 12, pp. 8076–8091, Dec. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9832954/>
- [25] Martin Ester, M. Ester, Hans-Peter Kriegel, H.-P. Kriegel, Jörg Sander, J. Sander, Xiaowei Xu, and X. Xu, "A density-based algorithm for discovering clusters a density-based algorithm for discovering clusters in large spatial databases with noise," *Knowledge Discovery and Data Mining*, pp. 226–231, Aug. 1996, mag ID: 2562836854 S2ID: 5c8fe9a0412a078e30eb7e5eeb0068655b673e86.
- [26] M. Ester and R. Wittmann, "Incremental generalization for mining in a data warehousing environment," in *Advances in Database Technology — EDBT'98*, G. Goos, J. Hartmanis, J. Van Leeuwen, H.-J. Schek, G. Alonso, F. Saltor, and I. Ramos, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, vol. 1377, pp. 135–149, series Title: Lecture Notes in Computer Science. [Online]. Available: <http://link.springer.com/10.1007/BFb0100982>
- [27] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 114, no. 13, pp. 3521–3526, Mar. 2017. [Online]. Available: <https://pnas.org/doi/full/10.1073/pnas.1611835114>
- [28] N. L. M. Jailani, J. K. Dhanasegaran, G. Alkaws, A. A. Alkahtani, C. C. Phing, Y. Baashar, L. F. Capretz, A. Q. Al-Shetwi, and S. K. Tiong, "Investigating the Power of LSTM-Based Models in Solar Energy Forecasting," *Processes*, vol. 11, no. 5, p. 1382, May 2023. [Online]. Available: <https://www.mdpi.com/2227-9717/11/5/1382>
- [29] G. Fungtammasan and I. Koprinska, "Convolutional and LSTM Neural Networks for Solar Power Forecasting," in *2023 International Joint Conference on Neural Networks (IJCNN)*. Gold Coast, Australia: IEEE, Jun. 2023, pp. 1–7. [Online]. Available: <https://ieeexplore.ieee.org/document/10191813/>
- [30] M. Abdel-Nasser and K. Mahmoud, "Accurate photovoltaic power forecasting models using deep LSTM-RNN," *Neural Comput & Applic*, vol. 31, no. 7, pp. 2727–2740, Jul. 2019. [Online]. Available: <http://link.springer.com/10.1007/s00521-017-3225-z>
- [31] V. Suresh, F. Aksan, P. Janik, T. Sikorski, and B. S. Revathi, "Probabilistic LSTM-Autoencoder Based Hour-Ahead Solar Power Forecasting Model for Intra-Day Electricity Market Participation: A Polish Case Study," *IEEE Access*, vol. 10, pp. 110628–110638, 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9921296/>
- [32] Meghana Bharadwaj and Sanjana Sarda, "Energy Prediction using Federated Learning," *arXiv.org*, Jan. 2023, arXIV_ID: 2301.09165 MAG ID: 4317940348 S2ID: be20ab605bacc726d85b3a6b3b56230a7e7b077a.
- [33] H. Wen, Y. Du, E. G. Lim, H. Wen, K. Yan, X. Li, and L. Jiang, "A solar forecasting framework based on federated learning and distributed computing," *Building and Environment*, vol. 225, p. 109556, Nov. 2022. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0360132322007867>
- [34] M. Savi and F. Olivadese, "Short-Term Energy Consumption Forecasting at the Edge: A Federated Learning Approach," *IEEE Access*, vol. 9, pp. 95949–95969, 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9469923/>