ACCURATE FEDERATED LEARNING WITH UNCERTAINTY QUANTIFICATION FOR DISTRIBUTED ENERGY RESOURCE FORECASTING APPLIED TO SMART GRIDS PLANNING AND OPERATION: THE ALAMO VISION

Lucas Pereira[1,2*], Vineet Nair[3], Anuradha Annaswamy[3], Bruno Dias[4], Hugo Morais[1,5]

[1]Instituto Superior Técnico - IST, University of Lisbon, Lisbon, Portugal
[2]Interactive Technologies Institute, LARSyS, Lisbon, Portugal
[3] Massachusetts Institute of Technology, Cambridge, MA, USA
[4]Universidade Federal de Juiz de Fora, Juiz de Fora, MG, Brazil
[4]INESC-ID – Instituto de Engenharia de Sistemas e Computadores - Investigação e Desenvolvimento, Lisbon, Portugal
**lucas.pereira@tecnico.ulisboa.pt*

Keywords: FEDERATED LEARNING, UNCERTAINTY, DER, PLANNING AND OPERATION.

Abstract

This paper presents the ALAMO vision, which is a research project that aims at developing technologies to facilitate the planning and operation of power grids with very high penetration of Distributed Energy Resources while at the same time assuring the privacy of the main actors on the demand side. The ALAMO project will address outstanding research challenges related to the development of accurate forecasting algorithms based on Federated Learning (e.g., forecasting in-front and behind-the-meter PV production), as well as challenges associated with producing sharped and well-calibrated quantifications of epistemic and aleatoric uncertainty for such forecasting models. Finally, use cases will be carefully crafted to understand how to incorporate FL forecasts and uncertainty into operational planning and operation tools. The use cases will be demonstrated in virtual and physical testbeds in Portugal, USA, and Brazil.

1. Introduction

As Renewable Energy Sources (RES) begin to dominate the generation portfolios, it becomes crucial to produce accurate forecasts and quantify the uncertainty of the developed forecasting models to assess their predictive confidence, which is critical for the real-world decision process. Uncertainties can impact several tasks needed in power systems, such as security assessment, operational planning, wholesale electricity markets, or hosting capacity. The impact on these processes can be the violation of technical operating constraints, non-compliance with market operations, or imposing unnecessary limits on hosting capacity [1].

The main objective of the Accurate federated Learning with uncertainty quantification for DER forecasting Applied to sMart Grids planning and Operation (ALAMO[1]) project is to develop technologies to facilitate the management of power grids with very high penetration of Distributed Energy Resources (DERs), while at the same time assuring the privacy of the main stakeholders. In particular, the project will address outstanding research challenges related to developing accurate forecasting algorithms based on Federated Learning (FL) and challenges associated with producing sharp and well-calibrated quantifications of epistemic and aleatoric uncertainty for such forecasting models.

Federated learning is a decentralized approach that enables collaborative training of Machine Learning (ML) models in distributed environments where data is stored in silos, i.e., locally on different devices or systems within an organization. This approach addresses privacy concerns, as raw data is not directly transmitted to a centralized location, and improves the efficiency of model training in distributed environments by reducing storage and communication costs, which differentiates it from the centralized ML environment [2], [3], [4].

Predictive uncertainty can be categorized into two classes: 1) model (epistemic) and 2) data (aleatory) uncertainty [5] [13]. The former indicates how uncertain the model is when explaining the observed data. The latter describes the variance of the conditional distribution of the target variable given input features. This uncertainty arises from the stochastic nature of the observed input features, such as noise or measurement errors.

We posit that accurate forecasting models can be built using FL even though the accuracy is still far from the results

[1]ALAMO Project: https://alamo.tecnico.ulisboa.pt

obtained using traditional centralized forecasting models [6]. Similarly, we postulate that Sharp and well-calibrated quantifications of uncertainties are possible in FL-based methods. Quantifying uncertainty is an area with very few works, particularly in the case of FL, where only a few existing works have attempted to extend methods developed for centralized forecasting [7].

The developed forecasting and uncertainty models will be incorporated into operational planning tools for distribution system operators (DSOs). Such tools are becoming crucial for DSOs due to increased RES connected to the distribution systems. Yet, simultaneously, the same tools are quickly becoming obsolete since they need to consider the consumers' privacy or uncertainty in the employed forecasts. i.e., current tools assume full observability of the grid and perfect forecasting in different horizons.

This paper details the proposed research framework and briefly describes two datasets that will be used in the project and four FL frameworks that have been studied.

2. ALAMO Research Framework

The ALAMO project consists of five tasks:

1. Dataset Preparation and Grid Modelling,
2. Development of FL Forecasting Algorithms,
3. Uncertainty Quantification and Scoring,
4. Application of FL Forecasts for Grid Operational Planning, TSO-DSO Coordination, and P2P Trading,
5. Demonstration in Real-World Testbed.

Each of the objectives and the research plan is briefly described next.

2.1. Dataset Preparation and Grid Modelling

Power distribution systems are changing rapidly with the increasing penetration of flexible loads and other DERs. This task focuses on dataset collection, cleaning, and curation, as well as employing accurate grid models that better represent the future paradigm of distribution systems and their impacts on the transmission system and power grids. There will also be an emphasis on developing open-source data repositories and associated grid models.

2.2. Development of Federated Learning Forecasting Algorithms

Prior works have looked at applying Federated Learning (FL) for load and Distributed Energy Resources (DER) forecasting [8], but there are still several limitations and challenges remaining that need to be addressed. These include (i) poorer performance of FL in terms of prediction accuracy when compared to traditional centralized learning methods, mainly when applied to data that is not independently and identically distributed (IID) across different FL clients, and (ii) long training times for the aggregated global model to converge due to diverse weights and parameters across heterogeneous clients [9].

The ALAMO project aims to address these challenges by studying recent variations of FL, like hierarchical clustering (HC) [10], iterative federated clustering algorithm (IFCA) [11], and model personalization or fine-tuning [12]. Hybrid approaches will also be tested considering the combination of several methods.

2.3. Uncertainty Quantification and Scoring

Ongoing research efforts concerning the quantification of uncertainty in FL have focused on extending existing methods from centralized forecasting, such as ensembles and the Monte Carlo Dropout (MCD) [7].

The ALAMO project will instead explore methods based on Quantile Regression (QR) for uncertainty quantification. Compared to other approaches, e.g., Monte MCD, QR can model complex distribution without making any *apriori* assumptions on the underlying distribution of the data [13]. This project will propose a new scoring metric to quantify the predictive uncertainty of the forecasting models by exploring the trade-offs between sharpness, predictive calibration probability, and predictive forecasting error of the trained FL models. Finally, the project will also assess the generalization potential by conducting sensitive studies for various forecasting horizons and temporal resolutions and evaluate the efficacy and quality of the uncertainty estimates using Auto-Regressive Inference (ARI).

2.4. Application of FL Forecasting for Grid Planning

As additional sources of uncertainty appear in the power systems, it becomes crucial to understand the impact of such uncertainties on system operators' and end-users' activities. The ALAMO project will evaluate the impact of uncertainties in three main processes: 1) operational planning tools; 2) Transmission System Operators (TSO) / Distribution System Operators (DSO) Coordination; and 3) peer-to-peer negotiation between small consumers and/or prosumers.

Concerning the effects of uncertainty in operational planning tools, the main aim is to define the system's hosting capacity one day ahead (or in another period) and evaluate the best strategies for activating flexibility. Uncertainties will be introduced in the existing algorithms, such as the optimization based on a flexibility sensibility matrix or based on a stochastic AC-OPF. Concerning the TSO/DSO coordination, a new method with support for uncertainties will be proposed. Three services will be tested: frequency regulation, volt-var control, and congestion management. In those services, the requirements can be requested by both TSOs and DSOs. However, instead of deterministic information exchange, probabilistic information will be exchanged. Finally, a P2P model based on bilateral trades will be developed considering uncertainties in production (selling offers) and consumption (buying offers). In this model, prosumers will trade with each other, deciding the amount of power exchanged and the corresponding trading price.

2.5. Demonstration

Most research in FL is being developed in sandboxes without considering real-world constraints. In this regard, the ALAMO project aims to demonstrate the developed forecasting algorithms in real-world settings leveraging existing real-world deployments of smart meters. Each site is equipped with a gateway based on Raspberry Pi, and they are all connected to a cloud-based server that integrates the data every minute. The developed algorithms will be deployed in the edge devices, and the performance of FL will be benchmarked (in real-time) with that of the centralized models.

3. Preliminary Results

3.1. Datasets

So far, two datasets have been identified: i) the Madeira Island Prosumer Dataset and ii) the Portugal Mainland Secondary Substation Dataset.

The Madeira Island Prosumer dataset was collected in the scope of the Horizon 2020 Smart Island Energy Systems (SMILE) project and consists of aggregated measurements for electric energy consumption and production from the PV [14]. The dataset comprises a mix of households and small businesses, with 27 participants. The data was collected between 2018 and 2021, comprising around three years of measurements. The original sampling rate of the data is 1Hz and contains the following measurements: current, voltage, active, reactive, apparent power, and frequency. Finally, the dataset was supplemented with solar radiation and temperature data taken from Solcast [15] during the same period – using the approximate coordinates of each prosumer. This data is available at 5, 10, 15, 30 and 60 minutes.

The Portugal Mainland Secondary Substation dataset contains data measured from 179 secondary substation datasets – active and reactive power. From these, 62 have measurements for production and consumption, while the remaining 131 contain only consumption measurements. The data is supplemented with meteorological data from two Portuguese systems: IPMA[2] (Instituto Português do Mar e da Atmosfera) and IST Meteo[3]. This dataset's granularity is 1 sample every hour, and measurements are available between January 1st, 2019, and December 31st, 2020.

3.2. Federated Learning Frameworks

One of the project's first steps was to identify existing open-source FL development frameworks through a desk survey. Ultimately, it was possible to find eight such frameworks, four of which were explored in more detail: NVIDIA FLARE[4], FATE[5], Flower[6], and TensorFlow Federated[7].

The NVIDIA FLARE platform supports deep learning and traditional machine learning algorithms from common frameworks like PyTorch, TensorFlow, and Scikit-learn. It facilitates horizontal and vertical FL, incorporating built-in algorithms such as FedAvg, FedProx, FedOpt, Scaffold, and Ditto. FLARE supports various training and validation workflows and offers data analytics and machine learning lifecycle management features. Privacy preservation methods like differential privacy, homomorphic encryption, and private set intersection are also integrated. FLARE has been successfully applied in the healthcare sector. Its pros include suitability for commercial use, framework-agnostic nature, support for tree-based models, and tools for monitoring and provisioning. However, it may require some time since it has a steep learning curve.

The FATE (Federated AI Technology Enabler) framework offers federated machine learning algorithms covering most parts of the ML pipeline: Data Input/Output, data preprocessing, feature engineering, and modeling. It includes utilities for encryption, statistics, parameter definitions, and auto-generation of transfer variables. FATE provides kits and base models for creating new algorithm modules with reusable functions for standardization and compactness. It offers multiple security protocols for secure multi-party interactions. FATE has been utilized for credit risk control, object detection, and anti-money laundering tasks. Its pros include suitability for commercial use, a wide range of federated learning algorithms, additional features, and compatibility with Spark clusters. However, users may need help to extend and be more user-friendly.

The Flower, originating from a research project at the University of Oxford, is a highly flexible FL framework designed to cater to individual use cases. It supports various configurations and can be extended extensively, making it suitable for diverse AI research needs. It offers ease of use and framework-agnostic flexibility and is compatible with multiple machine learning frameworks like PyTorch, TensorFlow, and sci-kit-learn. Written with maintainability in mind, Flower encourages community contributions. It has been applied successfully in tasks like image classification. Its pros include ease of use, flexibility, framework-agnostic nature, and suitability for research problems. However, it may need more features than other frameworks.

TensorFlow Federated (TFF) is an open-source framework for federated learning consisting of two layers: Federated Learning (FL) and Federated Core (FC). FL provides high-level interfaces for plugging existing machine learning models into TFF without delving into the details of

[2] IPMA: https://www.ipma.pt/pt/index.html
[3] IST Meteo: https://meteo.tecnico.ulisboa.pt/
[4] FLARE: https://developer.nvidia.com/flare
[5] FATE: https://fate.fedai.org/overview/
[6] Flower: https://flower.ai/
[7] TFF: https://www.tensorflow.org/federated

federated learning algorithms. At the same time, FC offers lower-level interfaces for expressing custom federated algorithms using TensorFlow and distributed communication operators within a functional programming environment. TFF has been effectively used for tasks like image classification and text generation. Its pros include being well-suited for TensorFlow and Keras. However, this is also a limitation since it does not support other relevant frameworks.

4. Conclusion

This paper presented the vision of the ALAMO research project that aims to set itself as a framework to improve federated DER forecasting and assess the impact of forecasting errors and uncertainties in the context of smart grid planning and operational tools. Two datasets from Portugal were also presented, as well as an overview of four FL frameworks commonly used in the FL literature.

5. Acknowledgments

This work is partially funded by the Portuguese Foundation for Science and Technology (FCT) under a grant (2022.15771.MIT) through national funds. FCT also supports the authors under grants UIDB/50009/2020, UIDB/50021/2020 and CEECIND/01179/2017.

6. References

[1] A. Faustine and L. Pereira, ‘FPSeq2Q: Fully Parameterized Sequence to Quantile Regression for Net-Load Forecasting With Uncertainty Estimates’, *IEEE Transactions on Smart Grid*, vol. 13, no. 3, pp. 2440–2451, May 2022, doi: 10.1109/TSG.2022.3148699.

[2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, ‘Communication-Efficient Learning of Deep Networks from Decentralized Data’, in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, PMLR, Apr. 2017, pp. 1273–1282. Accessed: Mar. 05, 2024. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[3] A. Z. Tan, H. Yu, L. Cui, and Q. Yang, ‘Towards Personalized Federated Learning’, *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 12, pp. 9587–9603, Dec. 2023, doi: 10.1109/TNNLS.2022.3160699.

[4] P. Qi, D. Chiaro, A. Guzzo, M. Ianni, G. Fortino, and F. Piccialli, ‘Model aggregation techniques in federated learning: A comprehensive survey’, *Future Generation Computer Systems*, vol. 150, pp. 272–293, Jan. 2024, doi: 10.1016/j.future.2023.09.008.

[5] A. Loquercio, M. Segu, and D. Scaramuzza, ‘A General Framework for Uncertainty Estimation in Deep Learning’, *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3153–3160, Apr. 2020, doi: 10.1109/LRA.2020.2974682.

[6] E. Petrangeli, N. Tonellotto, and C. Vallati, ‘Performance Evaluation of Federated Learning for Residential Energy Forecasting’, *IoT*, vol. 3, no. 3, Art. no. 3, Sep. 2022, doi: 10.3390/iot3030021.

[7] F. Linsner, L. Adilova, S. Däubener, M. Kamp, and A. Fischer, ‘Approaches to Uncertainty Quantification in Federated Deep Learning’, in *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, vol. 1524, M. Kamp, I. Koprinska, A. Bibal, T. Bouadi, B. Frénay, L. Galárraga, J. Oramas, L. Adilova, G. Graça, et al., Eds., in Communications in Computer and Information Science, vol. 1524. , Cham: Springer International Publishing, 2021, pp. 128–145. doi: 10.1007/978-3-030-93736-2_12.

[8] V. Venkataramanan, S. Kaza, and A. M. Annaswamy, ‘DER Forecast Using Privacy-Preserving Federated Learning’, *IEEE Internet of Things Journal*, vol. 10, no. 3, pp. 2046–2055, Feb. 2023, doi: 10.1109/JIOT.2022.3157299.

[9] Q. Yang, Y. Liu, T. Chen, and Y. Tong, ‘Federated Machine Learning: Concept and Applications’, *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, p. 12:1-12:19, Jan. 2019, doi: 10.1145/3298981.

[10] C. Briggs, Z. Fan, and P. Andras, ‘Federated learning with hierarchical clustering of local updates to improve training on non-IID data’, in *2020 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2020, pp. 1–9. doi: 10.1109/IJCNN48605.2020.9207469.

[11] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, ‘An Efficient Framework for Clustered Federated Learning’, in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2020, pp. 19586–19597. Accessed: Mar. 05, 2024. [Online]. Available: https://proceedings.neurips.cc/paper/2020/hash/e32cc80bf07915058ce90722ee17bb71-Abstract.html

[12] C. Briggs, Z. Fan, and P. Andras, ‘Federated Learning for Short-Term Residential Load Forecasting’, *IEEE Open J. Power Energy*, vol. 9, pp. 573–583, 2022, doi: 10.1109/OAJPE.2022.3206220.

[13] D. Kitahara, K. Leng, Y. Tezuka, and A. Hirabayashi, ‘Simultaneous Spline Quantile Regression Under Shape Constraints’, in *2020 28th European Signal Processing Conference (EUSIPCO)*, Jan. 2021, pp. 2423–2427. doi: 10.23919/Eusipco47968.2020.9287462.

[14] L. Pereira, J. Cavaleiro, and L. Barros, ‘Economic Assessment of Solar-Powered Residential Battery Energy Storage Systems: The Case of Madeira Island, Portugal’, *Applied Sciences*, vol. 10, no. 20, Art. no. 20, Jan. 2020, doi: 10.3390/app10207366.

[15] Solcast, ‘Solcast - Solar Forecasting & Solar Irradiance Data’, Solcast - Solar Forecasting & Solar Irradiance Data. Accessed: May 18, 2021. [Online]. Available: https://solcast.com/