A Federated Learning Approach for Distributed Human Activity Recognition

Article

Accepted version

F. Concone, C. Ferdico, G. Lo Re, M. Morana

In Proceedings of the 2022 IEEE International Conference on Smart Computing (SMARTCOMP), 2022, pp. 269-274,
doi: 10.1109/SMARTCOMP55677.2022.00066.

Publisher: IEEE

A Federated Learning Approach for Distributed Human Activity Recognition

Federico Concone, Cedric Ferdico, Giuseppe Lo Re, and Marco Morana
Università degli Studi di Palermo, Dipartimento di Ingegneria
Viale delle Scienze, ed. 6, 90128 Palermo, Italy
{federico.concone, cedric.ferdico, giuseppe.lore, marco.morana}@unipa.it

***Abstract*—In recent years, the widespread diffusion of smart pervasive devices able to provide AI-based services has encouraged research in the definition of new distributed learning paradigms. Federated Learning (FL) is one of the most recent approaches which allows devices to collaborate to train AI-based models, whereas guarantying privacy and lower communication costs. Although different studies on FL have been conducted, a general and modular architecture capable of performing well in different scenarios is still missing. Following this direction, this paper proposes a general FL framework whose validity is assessed by considering a distributed activity recognition scenario in which users' personal devices are employed as the basis of the sensing infrastructure. Experimental analysis was performed to evaluate the effectiveness of the architecture as compared with a centralized approach, under different settings. Results demonstrate the versatility and functionality of the proposed solution.**

***Index Terms*—Federated Learning, Distributed Computing, Machine Learning, Human Activity Recognition**

I. Introduction

Over the past decade, the diffusion of smart devices has driven the design of new Artificial Intelligence (AI) and Machine Learning (ML) solutions that require more and more computing power. Although simple machine learning models can be trained with modest amounts of data, more complex applications may require up to terabytes, or petabytes of training data. This results into the need to distribute the AI workload across multiple machines over the network.

Distributed Learning (DL) was introduced to overcome the problem raised on both the computational and storage sides of heavy ML models. The core idea of DL is to distribute the computation along a cluster of devices capable to perform part of the learning process, while a single *central entity* is responsible for aggregating the output from every participant so as to obtain a better unified model.

Due to its intrinsic nature, Distributed Human Activity Recognition (DHAR) represents the ideal scenario where DL can provide significant improvements to the system performance. The benefits brought by the distributed paradigm are countless, but protecting sensitive data is not one of them. In fact, given the central role of people in the sensing process, preserving sensitive information is mandatory and, generally, security mechanisms are designed ad-hoc to fulfill this requirement. This represents a critical aspect in the designing of a distributed architecture [1]. To better comprehend the issue, consider a healthcare application in which ML models are trained on data produced by user's wearable device to track health status, e.g., assess if a hospital's patient performs sufficient physical activities during the day. Here, the critical challenge is that all data must be stored and processed within the hospital. In real life, it is not possible to share user's sensitive data with other organizations due to privacy and security concerns, thus making difficult to train powerful models because of the lack of required amount of data [2].

Federated Learning (FL) is a very specific distributed machine learning paradigm that differs from others in at least three key points [3]. The first is that, in a FL scenario, direct raw data communications among parties are not allowed. Since raw data may have multiple ownerships, this makes FL approaches intrinsically compliant with the related laws. Moreover, FL allows to exploit the distributed computing resources in multiple regions or organizations, rather than a single server, or a cluster in a single region, belonging to a single organization. This is crucial to enable the collaboration among multiple organizations [4]. Finally, additional security mechanisms can be added if required by the application scenario [5]. Following these motivations, most of the related work is focused on exploiting the advantages of FL to realise "federated versions" of common ML algorithms, each according to its own architecture.

This paper presents a modular architecture for FL in which a set of *federated aggregators* play an intermediary role between the more highly distributed nodes and the centralised layer. Although such a solution could be employed in a wide range of scenarios just modulating its functional components, we chose to instantiate it to implement a DHAR scenario in which it is mandatory to protect sensitive data, i.e., the activities performed by users. In particular, wearable devices aim to collect data from embedded sensors, and share them with the federated aggregators at the upper layer. Then, the federated aggregator uses raw data to (i) infer the activity performed by the user and (ii) update local ML models to be shared with the central entity, the only one responsible for the updating of the global model. We evaluate performance of our architecture in such a composite scenario and demonstrate that federated learning is a good trade-off between privacy, performance and versatility compared to a centralized learning approach for the training of the HAR classifier.

The remainder of the paper is organised as follows: Related work is outlined in Section II. The architecture and its main

modules are described in Section III. Experimental setup and results are presented in Section IV. Conclusions will follow in Section V.

II. Related Works

In recent years, research on Human Activity Recognition (HAR) is gaining significant attention since several application scenarios exist in which HAR can be successfully adopted, e.g., urban mobility management [6], ambient intelligence [7], [8], and assisted living [9].

HAR is generally achieved by exploiting two different types of input data, namely visual or sensory.

The common idea behind vision-based HAR is to describe users by means of silhouettes that allow to extract features about their movements, from RGB or depth images, and then perform the activity recognition trough a machine learning model [10]. Several applications have been proposed following this methodology, especially in the context of intelligent environmental systems [11], such as energy management [12], or smart surveillance [13]. Despite the benefits, vision-based HAR is quite limited to indoor environments and particularly heavy in terms of computational burden.

Sensor-based HAR techniques have been then proposed as a possible alternative. In this context, human activities can be intuitively considered as sequences of recurrent patterns in raw data captured from sensors worn by the users, such as those attached to wearable elastic bands or embedded in smartphones, smartwatches, etc. Many HAR algorithms have been presented in the literature and, generally, differ based on the type (and number) of sensors used, or the features extracted from the sensor readings. For example, in [14] authors present a framework for HAR using data captured by means of triaxial accelerometer and gyroscope sensors embedded into the smartphone. The temporal patterns generated by these sensors are firstly analysed to model activities via a 30-dimensional feature vector, and then classified according a machine learning approach. Another interesting example presented in literature is the one implemented in the Google Activity Recognition APIs for Android [1]. However, this tool acts as a black box by not providing a way to understand what features are being used, nor the model used for classification.

All of the approaches discussed above suffer from two issues stemming from (i) the growing demand for increasingly powerful smart models and services, and (ii) user concerns about sharing sensitive data (e.g., the activity performed) with third parties.

Over the years, the first issue was alleviated by proposing scalable and time-efficient solutions that exploit *Cloud Computing*, *Fog Computing*, or *Edge Computing* paradigms to provide HAR services. While *Cloud Computing* could provide a feasible solution to move heavy computation towards the cloud, its applicability in real-time applications is negligible as data is continuously transferred from/to the cloud. Then, the remaining *Fog Computing* and *Edge Computing* have been strongly investigated in HAR scenarios. In [15], a fog-based architecture for complex human activity recognition is proposed. Here, sensor readings are processed as close as possible to data source so that it is possible to meet real-time constraints. In particular, the machine learning model runs on powerful entities within the network (i.e., Fog entities) because of the complexity of the activity recognition, i.e., a combination of K-means clustering, Support Vector Machines, and Hidden Markov Models.

The idea of distributing heavy computations among intermediate (fog) or remote (cloud) devices alleviates for sure the growing demand for more powerful services, but it does not address privacy or security issues. In fact, because both of their distributed nature and high degree of modularity, edge-fog-cloud computing systems are particularly prone to cyber security attacks that can be performed against every element of the infrastructure [16].

In order to meet the requirements of distributed computing and privacy preservation, Federated Learning (FL) has been proposed in 2016 [17]. It enables a large number of edge devices, called *clients*, storing local data observations to locally and collaboratively train one single machine learning model without having to share their raw data. A coordinating server then aggregates the contributions from all the edge devices and shares an updated model with the participating clients to benefit from their learning experience.

Several applications have benefited from the FL paradigm, ranging from finance [18] and monitoring [19], to healthcare [20] scenarios. Often the healthcare scenario itself heavily leverages an Activity Recognition process. For this reason, HAR is assuming a role of great interest to the research community in the federated learning domain. In [21], the authors have evaluated a Softmax regression and a deep neural network for the task of HAR. The results demonstrated the improvement the FL may bring, allowing to achieve acceptable accuracy while preserving privacy. In [2], a federated transfer learning framework, called FedHealth, for wearable healthcare is discussed. Through federated learning and homomorphic encryption, FedHealth aggregates the data from different organizations to build powerful machine learning models with the users' privacy well preserved. After the cloud model is built, FedHealth utilizes transfer learning methods to achieve personalized model learning for each organization.

According to the literature, Federated Learning exhibits clear theoretical advantages over classical centralized learning from a pervasive computing perspective. But little is known about how these advantages are actually achieved in practice, and the behavior of such learning approaches [22]. Motivated by these reasons, we propose a FL architecture to evaluate how it may work on the HAR task.

III. Proposed Architecture

The modular architecture we propose is deployed on the hierarchical topology depicted in Fig. 1.

At the lowest level, sensing devices (SDs) are responsible for collecting raw data and, if required, performing simple data

[1] https://developers.google.com/location-context/activity-recognition

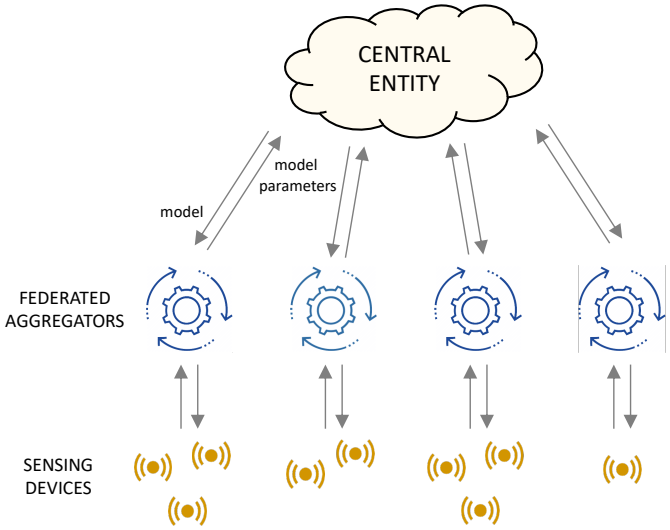

Fig. 1: Federated Learning Architecture.

pre-processing. Then, all the data is shared with the entities at the upper layer through classical network protocols. Since the system processes rough data within the federation, information at this layer is not encrypted thus meeting the computing constraints of the adopted smart devices, if any.

At the intermediate level, *Federated Aggregators* (FAs) are designed to perform in-depth analysis on data obtained previously. In particular, the main goal of an FA is to refine its inner model, extract local model updates, and share this information to the upper layer. Other equally important FA's tasks are the collection/management of new data, through which the model is updated, and the actual classification exploited for the service provisioning. The entities at this level could be companies that want to collaborate to achieve a common goal, without sharing their sensitive data.

Information produced by FAs is sent to the *Central Entity* (CE) which is responsible for aggregating the local models' parameters coming from all the underlying devices, and updating the global model by means of its internal parameters. The results of this analysis are sent back to the FA in order to update their behaviours, making the whole system consistent.

A. Modules for Federated Leaning

The main modules that constitute the *federated aggregators* and the *central entity* are detailed in Fig. 2.

The CE represents the heart of the whole architecture, since its main role is to coordinate each FA. In our general view, this entity is the service provider and is responsible for performing a preliminary training phase on the data contained in the *Start Dataset*. After training, the *Federated Classifier* is tuned to fulfill the application specific goals, and the resulting (internal) parameters are shared with the underlying FAs through the *Model Distributor*.

Each FA receives the global model via the *Model Receiver*, and forwards this information to its own *Federated Classifier*. The *Model Receiver* may also pass the model obtained from CE to the *Model Sender*; this is useful for that kind of applications that require to calculate some parameters, e.g., *crossover point* or *regret* [23], to evaluate the participation at the federated learning.

The classifier used in FA is the same as the one held by the CE, but it is the one that actually performs the classification. The prediction then is passed as input to the *Actuator*, which is responsible for moving and controlling a mechanism, for example by activating a service or a functionality in the sensing device.

It is important to note that FAs are also expected to support the update of the global model located in the CE. In this regard, the *Data Manager* firstly collects new raw data from the sensing device and, then, extracts the features before sending them to the *Federated Classifier* and the *Local Dataset* for prediction and storage tasks, respectively.

Finally, apart from the *Model Sender* and *Trainer* modules which operate similarly to their counterparts in the CE, each FA is equipped with an additional module called *Local Update Manager*. This is triggered every time a certain condition is satisfied, e.g., a considerable amount of new data has been collected, or a certain amount of time has passed.

B. Human Activity Recognition

A straightforward approach in the federated scenario could be to exploit users' personal devices, such as smartwatches and smartphones, to act as sensing devices. These are able to (i) collect raw data from embedded sensors while user performs a particular activity, and (ii) make preliminary elaborations if they are necessary. A single FA with a higher level of performance, e.g., a personal computer, can be used to process and aggregate data from multiple devices worn by a community of users. We want to put in evidence that the system may use information directly captured by users' smartphones. These devices would be logically located in the bottom layer of the architecture, whilst the FA could consist of other types of units. More generally, given the proposed architecture, FAs can be any device with enough computing power to perform raw data analysis and aggregation, such as smart lighting poles, bus shelters, or any other structure deployed in a smart city.

For instance, we could imagine a Central Entity interested in offering smart services for automatic customization of SD settings. The sensing device may be set on silencing mode if the FA detects that the user is *relaxing* or *sleeping*. In such a composite scenario, we can imagine the CE to be the service provider, the FA is the user's smartphone, and the SD is represented by less powerful devices such as wrist-worn or smartwatches. Similarly, we can consider a federated video surveillance scenario in which cameras can be used for different purposes, e.g., recognising suspicious activity and sending the images to some local processing unit responsible for maintaining AI models. In this case, SDs are the cameras located in the environment, the FAs may be devices at disposition of the public entities, such as servers or laptops, while the CE could be the provider that offers the suspicious activity recognition service.

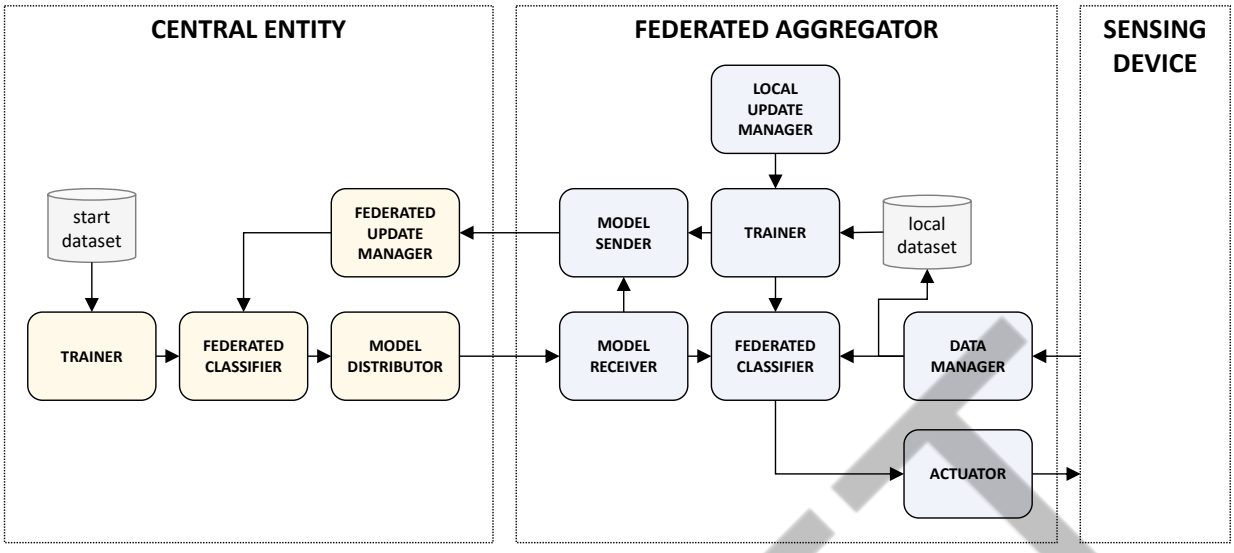

Fig. 2: Proposed FL modular architecture.

In the addressed HAR scenario, users may wear smartwatches while performing their activities, passing heavy data processing to more powerful FA devices, e.g., the users' smartphones [15]. Activities to recognise can be traced by the accelerometer and gyroscope embedded in the smartwatches. The collected patterns are sent to the FA, and specifically to the *Data Manager* module, to extract feature vectors for each time window. The FA of our case study, i.e., the users' smartphones, performs activity recognition by means of a Deep Neural Network (DNN), properly implemented in the *Federated Classifier* module. Moreover, the *Federated Aggregator* stores local model parameters, and any other data which need to be transferred to the CE. The *Central Entity* receives and then aggregates the models' updates in the *Federated Update Manager* via the *FedAvg* algorithm [24]. Finally, the improved global model is broadcast back to the FAs thanks to the *Model Distributor*.

IV. Experimental Evaluation

All the experiments were conducted on a public dataset [25] that includes 6 activities (*Walking*, *Standing*, *Laying*, *Sitting*, *Walking Upstairs*, *Walking Downstairs*) performed by 30 volunteers aged 19-48. Each activity is described by means of accelerometer and gyroscope, with both sensors generating raw data at a constant rate of 50Hz. The signals extracted from the sensors were pre-processed by applying noise filters and sampled in fixed-width sliding windows of 2.56 sec and 50% overlap resulting on 128 readings per window. Then, each activity is modeled as a feature vector containing more than 500 values in time and frequency domains.

In view of using these data in the Federated Learning scenario, we divided the full dataset in partitions such that each contains all the activities carried out by a single user and presents almost the same number of samples. In this way, we assigned each sensing device its own local dataset to which new data could be added later, thus simulating runtime behavior. Moreover, data regarding three users were kept out of the learning pool for testing purposes.

TABLE I: DNN network configuration.

DNN model parameters	
No. of layers	*6*
Activation function	*ReLU, SoftMax* (output layer)
Optimizer	*SGD*
Learning rate	*0.001*
Momentum	*0.8*
Training Epochs	*20*
Batch size	*32*

The metrics we employed for the evaluation are *Accuracy*, *Precision*, and *Recall*. Since this work deals with a multi-class classification problem, the final score of each metric is measured as the mean for each class.

Two different types of experiments were conducted to compare the performance of our federated solution with a centralized architecture. In particular, we refer to the first scenario as *static*, since the number of the sensing devices is the same for all the FAs, while the second scenario simulate a *dynamic* environment in which the number of SDs associated with each FA change over the time. In both settings, we considered three FAs that represents a good trade-off in terms of diversity and efficiency.

Firstly, we looked for the best set of parameters for the activity recognition procedure, that is the best hyperparameters for the DNN. To achieve this goal, a 10-fold cross validation was performed and result in the final configuration shown in Table I.

Then, regarding the *static* scenario, the tests were conducted by simulating a CE whose initial model is trained on its local data, and subsequently refined by parameters received from FAs. The results of this analysis are summarized in Fig. 3. Here, the baseline represents the centralized model and

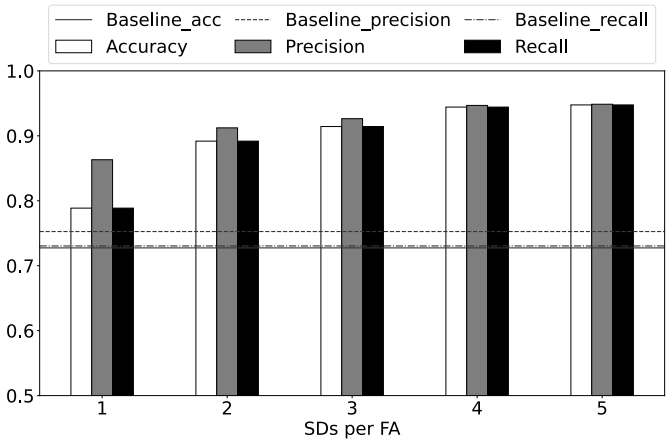

Fig. 3: Aggregated model's Accuracy, Precision and Recall test results in function of the number of users. The distribution of users per Federated Aggregator is homogeneous.

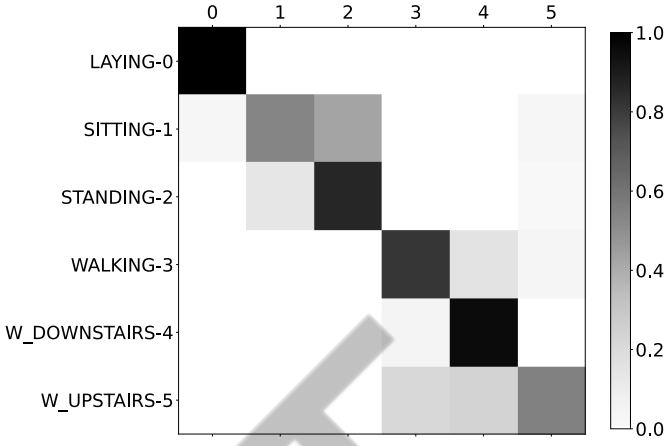

Fig. 4: Confusion matrix obtained by considering only one sensing device for each FA.

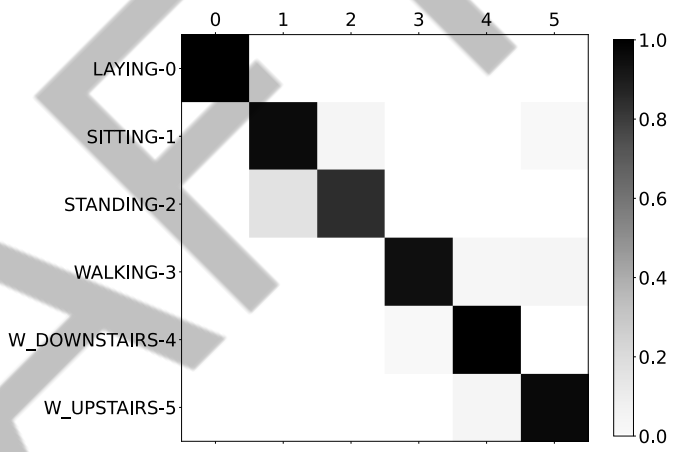

Fig. 5: Confusion matrix obtained by considering five sensing device for each FA.

is depicted as a constant (for all the metrics) since sensing devices do not share their local data (i.e., the central model is never updated). It is possible to observe an improvement in the performance of the federated model already from a small number of sensing devices. This is due to the fact that as data availability increases, the model manages to improve its performance even without ever having direct access to them.

To better understand the reasons for this marked difference, we present the confusion matrices obtained by the system when the number of SDs for each FA is 1 (Fig. 4) and 5 (Fig. 5). In particular, each $C_{i,j}$ cell represents the number of occurrences in class i that have been classified by the system as belonging to class j. Darker cells correspond to higher values, up to a maximum of 1. Therefore, main diagonal values correspond to true positives, and values outside the diagonal indicate classification errors. Ideally, we would like to get a very dark main diagonal, and lighter values in the other cells, which would indicate a low degree of confusion between activities.

Fig. 4 shows that the system relying on only one SD per FA has difficulty in correctly discriminating the activities; for instance, *Sitting* is often confused with *Standing*, and *Walking* with *Walking Downstairs*. Moreover, with this specific configuration, *Walking Upstairs* is strongly misclassified against *Walking* and *Walking Downstairs*. These difficulties can be easily explained by the fact that information available on the CE is not sufficient to well describe these activities. As expected, increasing the number of sensing devices for the FAs overcomes the problem. The confusion matrix in Fig. 5 is characterized by a very marked main diagonal, which shows how the system is able to recognise all activities satisfactorily, without confusing them with each other.

Other experiments focused on the *dynamic* case where the system was evaluated over the time. In particular, for each timestamp a number of sensing devices (in range $[1, 3]$) was added, and randomly distributed in one of the three FAs. As it is possible to observe in Fig. 6, even in this case the results were highly acceptable already from the first timestamps in which the number of SDs was low. Nevertheless, the performance of the system, in its federated version, grows as the number of the sensing devices increase (and therefore in the availability of data), presenting satisfactory performance from quite low numbers, exceeding the threshold of 95% for all metrics considered as the number of SDs grows.

V. Conclusions

In this paper, we presented an architecture for recognizing human activities through users' smart devices. The recognition process relies on a Federated Learning architecture where *Sensing Devices*, *Federated Aggregators*, and a *Central Entity* cooperate, at three different logic layers, for collecting sensory data, performing HAR, and improving the overall model.

The experiments were performed on a public dataset and aim to compare the performance of the proposed solution against a centralized approach. The outcomes prove that the FL architecture is able to generate an adequate model for the

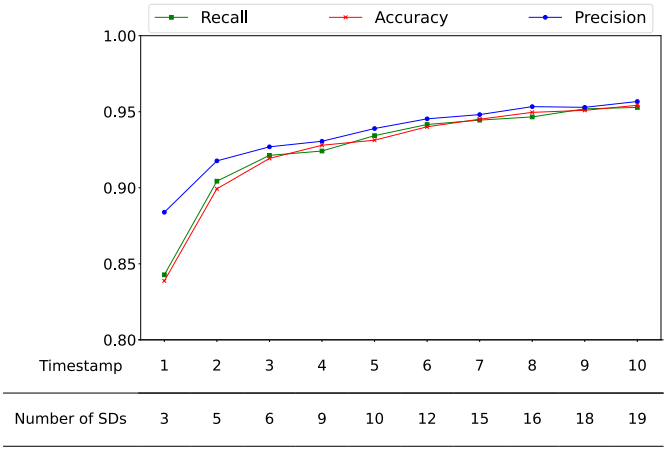

Fig. 6: Aggregated model's Accuracy, Precision and Recall trends in function of the number of users for each timestamp. The distribution of users per Federated Aggregator is not homogeneous.

HAR task, while the centralized version continues to achieve not satisfactory performance. Our solution constitutes a good trade-off between privacy, performance and versatility.

As future work, we plan to design a reputation management mechanism [26] able to mitigate the presence of malicious entities aiming at compromising the global model by injecting (intentionally) erroneous data to the *Central Entity*.

VI. Acknowledgment

This research is partially funded by the Project CrowdSense (PO FESR Sicilia 2014/2020).

References

[1] "A lightweight middleware platform for distributed computing on wireless sensor networks," *Procedia Computer Science*, vol. 32, pp. 908–913, 2014, the 5th International Conference on Ambient Systems, Networks and Technologies (ANT-2014), the 4th International Conference on Sustainable Energy Information Technology (SEIT-2014).

[2] Y. Chen, X. Qin, J. Wang, C. Yu, and W. Gao, "Fedhealth: A federated transfer learning framework for wearable healthcare," *IEEE Intelligent Systems*, vol. 35, no. 4, pp. 83–93, 2020.

[3] J. Liu, J. Huang, Y. Zhou, X. Li, S. Ji, H. Xiong, and D. Dou, "From distributed machine learning to federated learning: a survey," *Knowledge and Information Systems*, vol. 64, no. 4, pp. 885–917, Apr 2022.

[4] I. Kholod, E. Yanaki, D. Fomichev, E. Shalugin, E. Novikova, E. Filippov, and M. Nordlund, "Open-source federated learning frameworks for IoT: A comparative review and analysis," *Sensors*, vol. 21, no. 1, 2021.

[5] H. Zhu, H. Zhang, and Y. Jin, "From federated learning to federated neural architecture search: a survey," *Complex & Intelligent Systems*, vol. 7, no. 2, pp. 639–657, Apr 2021.

[6] A. Bordonaro, F. Concone, A. De Paola, G. Lo Re, and S. K. Das, "Modeling efficient and effective communications in vanet through population protocols," in *2021 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2021, pp. 305–310.

[7] S. Gaglio, G. Lo Re, M. Morana, and C. Ruocco, "Smart assistance for students and people living in a campus," in *2019 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2019, pp. 132–137.

[8] A. De Paola, P. Ferraro, S. Gaglio, G. Lo Re, M. Morana, M. Ortolani, and D. Peri, "An ambient intelligence system for assisted living," in *2017 AEIT International Annual Conference*, 2017, pp. 1–6.

[9] N. Gupta, S. K. Gupta, R. K. Pathak, V. Jain, P. Rashidi, and J. S. Suri, "Human activity recognition in artificial intelligence framework: a narrative review," *Artificial Intelligence Review*, Jan 2022.

[10] S. Gaglio, G. Lo Re, and M. Morana, "Human activity recognition process using 3-d posture data," *IEEE Transactions on Human-Machine Systems*, vol. 45, no. 5, pp. 586–597, 2015.

[11] F. Concone, P. Ferraro, and G. Lo Re, "Towards a smart campus through participatory sensing," in *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2018, pp. 393–398.

[12] A. De Paola, P. Ferraro, G. Lo Re, M. Morana, and M. Ortolani, "A fog-based hybrid intelligent system for energy saving in smart buildings," *Journal of Ambient Intelligence and Humanized Computing*, vol. 11, no. 7, pp. 2793–2807, 2020.

[13] V. Agate, F. Concone, and P. Ferraro, "Wip: Smart services for an augmented campus," in *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2018, pp. 276–278.

[14] F. Concone, S. Gaglio, G. Lo Re, and M. Morana, "Smartphone data analysis for human activity recognition," in *AI*IA 2017 Advances in Artificial Intelligence*, F. Esposito, R. Basili, S. Ferilli, and F. A. Lisi, Eds. Cham: Springer International Publishing, 2017, pp. 58–71.

[15] F. Concone, G. Lo Re, and M. Morana, "A fog-based application for human activity recognition using personal smart devices," *ACM Trans. Internet Technol.*, vol. 19, no. 2, mar 2019.

[16] ——, "Smcp: a secure mobile crowdsensing protocol for fog-based applications," *Human-centric Computing and Information Sciences*, vol. 10, no. 1, p. 28, Jul 2020.

[17] 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*, ser. Proceedings of Machine Learning Research, A. Singh and J. Zhu, Eds., vol. 54. PMLR, 20–22 Apr 2017, pp. 1273–1282.

[18] A. Abdallah, M. A. Maarof, and A. Zainal, "Fraud detection system: A survey," *Journal of Network and Computer Applications*, vol. 68, pp. 90–113, 2016.

[19] V. Agate, F. Concone, and P. Ferraro, "A resilient smart architecture for road surface condition monitoring," in *Innovations in Smart Cities Applications Volume 5*, M. Ben Ahmed, A. A. Boudhir, İ. R. Karaș, V. Jain, and S. Mellouli, Eds. Cham: Springer International Publishing, 2022, pp. 199–209.

[20] J. Xu, B. S. Glicksberg, C. Su, P. Walker, J. Bian, and F. Wang, "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, vol. 5, no. 1, pp. 1–19, Mar 2021.

[21] K. Sozinov, V. Vlassov, and S. Girdzijauskas, "Human activity recognition using federated learning," in *2018 IEEE Intl Conf on Parallel Distributed Processing with Applications, Ubiquitous Computing Communications, Big Data Cloud Computing, Social Computing Networking, Sustainable Computing Communications (ISPA/IUCC/BDCloud/SocialCom/SustainCom)*, 2018, pp. 1103–1111.

[22] S. Ek, F. Portet, P. Lalanda, and G. Vega, "Evaluation of federated learning aggregation algorithms: Application to human activity recognition," in *Adjunct Proceedings of the 2020 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2020 ACM International Symposium on Wearable Computers*, ser. UbiComp-ISWC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 638–643.

[23] S. Zehtabian, S. Khodadadeh, L. Bölöni, and D. Turgut, "Privacy-preserving learning of human activity predictors in smart environments," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, 2021, pp. 1–10.

[24] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-efficient learning of deep networks from decentralized data," 2016.

[25] D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. Reyes, *A public domain dataset for human activity recognition using smartphones*, 2013, p. 437–442.

[26] V. Agate, A. D. Paola, G. Lo Re, and M. Morana, "A simulation software for the evaluation of vulnerabilities in reputation management systems," *ACM Trans. Comput. Syst.*, vol. 37, no. 1–4, jun 2021.