Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)
Demonstrations Track

Fedstellar: A Platform for Training Models in a Privacy-preserving and Decentralized Fashion

Enrique Tomás Martínez Beltrán[1] , **Pedro Miguel Sánchez Sánchez**[1] , **Sergio López Bernal**[1] , **Gérôme Bovet**[2] , **Manuel Gil Pérez**[1] , **Gregorio Martínez Pérez**[1] and **Alberto Huertas Celdrán**[3]

[1]Department of Information and Communications Engineering, University of Murcia, 30100, Spain
[2]Cyber-Defence Campus, Armasuisse Science and Technology, 3602 Thun, Switzerland
[3]Communication Systems Group, Department of Informatics (IFI), University of Zurich, 8050 Zürich, Switzerland
{enriquetomas, pedromiguel.sanchez, slopez, mgilperez, gregorio}@um.es,
gerome.bovet@armasuisse.ch, huertas@ifi.uzh.ch

Abstract

This paper presents Fedstellar, a platform for training decentralized Federated Learning (FL) models in heterogeneous topologies in terms of the number of federation participants and their connections. Fedstellar allows users to build custom topologies, enabling them to control the aggregation of model parameters in a decentralized manner. The platform offers a Web application for creating, managing, and connecting nodes to ensure data privacy and provides tools to measure, monitor, and analyze the performance of the nodes. The paper describes the functionalities of Fedstellar and its potential applications. To demonstrate the applicability of the platform, different use cases are presented in which decentralized, semi-decentralized, and centralized architectures are compared in terms of model performance, convergence time, and network overhead when collaboratively classifying hand-written digits using the MNIST dataset.

1 Introduction

The increasing demand for Machine Learning (ML) and Deep Learning (DL) applications has led to a growing need for training large-scale models using distributed systems. With the proliferation of data-generating devices, such as smartphones, and the increasing demand for data privacy and security, it has become increasingly difficult to train models effectively and efficiently without compromising data privacy and security [Barbieri *et al.*, 2022]. In many cases, the data are distributed among different devices, organizations, or data centers, making it difficult to access, process, and analyze the data in a centralized manner [Nguyen *et al.*, 2022].

In recent years, Federated Learning (FL) has emerged as a promising solution for training ML/DL models using data from multiple devices or centers while preserving data privacy [McMahan *et al.*, 2016]. FL operates by distributing the model training process across multiple devices, such as smartphones, laptops, or IoT devices, instead of centralizing the data in a single location. Each device trains a local version of the model using its data and periodically sends updates to other nodes. These nodes can aggregate these updates to create a new model. The process repeats until the model converges or a stopping criterion is met.

Despite its potential, current FL platforms present several limitations, requiring a participant to act as an aggregator server, which can pose a bottleneck and limit scalability [Beltrán *et al.*, 2022]. Moreover, in specific scenarios, it is not possible for a participant to act as the aggregator server. To address these limitations, some versions of FL have been proposed to work in a decentralized fashion, distributing the computational load among multiple participants and eliminating the need for a single aggregator node. However, existing frameworks, such as TensorFlow Federated (TFF) or Federated AI Technology Enabler (FATE) [Li *et al.*, 2020], only cover Centralized Federated Learning (CFL), lacking in the literature platforms that fully support Decentralized Federated Learning (DFL) and Semi-Decentralized Federated Learning (SDFL) architectures. These federation architectures have different trade-offs regarding data privacy, scalability, and communication complexity, making it challenging to adapt existing platforms to different use cases and application domains.

To overcome these limitations, this demo paper presents Fedstellar, a platform for training ML/DL models in a collaborative, privacy-preserving, and decentralized fashion (publicly available in [Beltrán *et al.*, 2023]). The platform is composed of two main components: Fedstellar Core and a Web application. The first one is responsible for generating and deploying federation topologies in simulated or physical devices, such as smartphones, laptops, and embedded devices. It also supports different federation architectures, such as DFL, SDFL, and CFL, which are suitable for different use cases and application domains. The platform also allows for configuring federation nodes with customized datasets, models, or roles and supports different aggregation mechanisms. The second component is the Web application, which provides the user interface for generating and monitoring federated networks. It permits users to create and configure node

topologies and set up the federation parameters. The suitability of Fedstellar has been evaluated in three use cases with a fix number of federation nodes, different federation architectures, the well-known MNIST dataset, and a simple Convolutional Neural Network (CNN). In each use case, Fedstellar records and compares the average model accuracy per node, the network usage, and the participant's resource utilization.

2 Fedstellar Platform

The Fedstellar platform comprises two components: Fedstellar Core and a Web application to manage the deployment of scenarios. A scenario refers to a specific configuration of nodes and tasks in an FL architecture. Each node in the federation is equipped with the Fedstellar Core component, which provides the necessary functionalities for training FL models in a privacy-preserving way. The Web application serves as a central hub for the deployment and management of FL scenarios. Figure 1 shows the two components and their interaction on the platform.

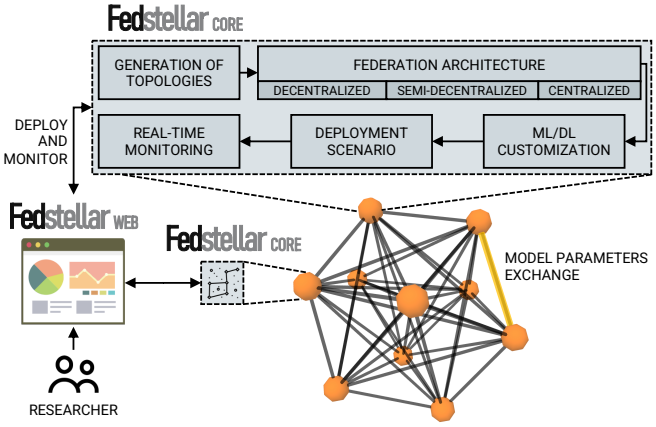

Figure 1: Integration of Fedstellar Core and Web application

Generation of Topologies. It provides the capability to generate different federated topologies, including fully connected and partially connected with star-structured, ring-structured, random, and node clustering configurations [Georgatos *et al.*, 2022]. The users can choose the topology that best fits their use case requirements. The fully connected network topology connects all nodes in the system, allowing for maximum collaboration and data exchange between all participants. The partially connected topologies allow for more efficient use of resources while still preserving privacy. Additionally, node clustering configurations can be used to model real-world scenarios where nodes may have limited connectivity or be grouped in clusters.

Federation Architecture. Fedstellar provides three different federation architectures: DFL, SDFL, and CFL [Beltrán *et al.*, 2022]. In DFL, all participants in the topology contribute to the FL process, and the model is updated locally at each node. In SDFL, some participants act as coordinators and aggregate the updates from other participants. In CFL, all model updates are sent to a central node, which aggregates the updates. Users can customize the federation scenario with aggregation algorithms, rounds, aggregation leadership, and other parameters [Hard *et al.*, 2021].

ML/DL Customization. Users can tailor the platform to their specific Artificial Intelligence (AI) requirements, including model and dataset characteristics. In this sense, it allows for using different algorithms and models, such as deep neural networks, for addressing classification, regression, or clustering tasks. Additionally, the platform can be adapted to different datasets, including structured and unstructured data, such as images, audio, and text.

Deployment Scenario. Fedstellar can deploy nodes on simulated environments, physical devices, and remote servers. It enables users to choose the deployment type best suits their needs and resources. Deployment on simulated environments provides fast and easy experimentation, while deployment on physical devices and remote servers allows real-world testing.

Real-time Monitoring. The platform incorporates geolocation and interactive graphs, providing valuable insights into performance and resource usage. The geolocation feature allows for the visualization of the spatial distribution of nodes. At the same time, the interactive graphs display model outputs in the training and evaluation phases and resource usage statistics, such as CPU or network traffic. This information can be used to optimize the federated process and make decisions regarding system configuration and deployment.

3 Demonstration

The Fedstellar platform underwent extensive evaluation through three use cases (UC1, UC2, UC3) to demonstrate its effectiveness in handling real-world scenarios. The evaluation involved a federation of ten participants represented as nodes in a network deployed in a fully connected topology using DFL (UC1) and SDFL (UC2). Finally, a star topology using CFL (UC3) was also tested to show its ability to handle centralization and decentralization in the network.

Figure 2a shows the user interface of the Web application providing the ability to select the scenario configuration through drop-down menus. In this demonstration, Fedstellar generated the UC1 scenario using a simulated deployment, a DFL federation architecture, and a fully connected topology, where all nodes were connected. The well-known MNIST dataset was selected for this demonstration as it provides a good benchmark for testing ML/DL platforms. The dataset was divided into not independent and identically distributed (non-IID) splits among the participants to simulate real-world scenarios where data is distributed among multiple parties. A straightforward CNN adapted to the data was utilized for training, and the FedAvg aggregation mechanism was used to combine the model parameters from each node in ten federation rounds. The same procedure was performed for UC2 and UC3 (see Figure 2), but it is not shown due to space reasons.

At this point, the nodes in the Fedstellar platform collaborated to create ML/DL models in a decentralized manner. Since each node creates collaborative AI models asynchronously in decentralized architectures, each node has a

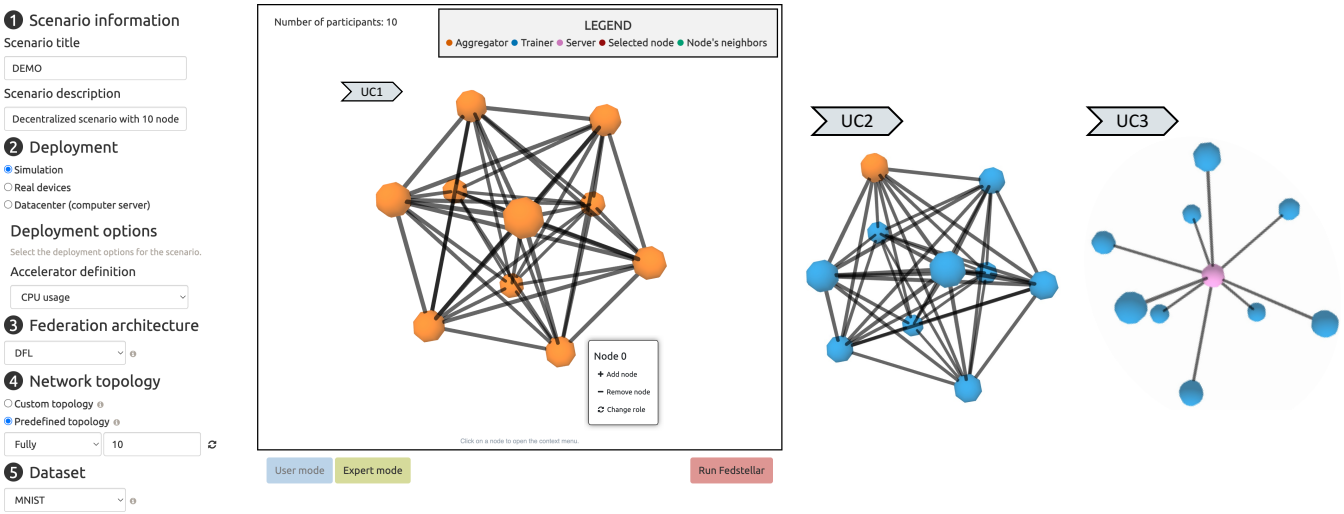

(a) Functionality of UC1 deployment using the proposed platform (b) Network topology deployed for UC2 and UC3

Figure 2: Setting up the Fedstellar platform for demonstration purposes

Use Case	Model (acc.)	Network (MB)	CPU (%)	Training Time* (s)
UC1 (DFL)	0.965	181	65	42
UC2 (SDFL)	0.955	163	62	53
UC3 (CFL)	0.967	132	41	57

* Overall time to reach model accuracy $\geq 90\%$

Table 1: Outcomes of Fedstellar platform. Runtime: 5 minutes

Figure 3: Network usage in MB based on federation architecture

different model from the other nodes. Performance metrics were then collected from each node with an emphasis on metrics related to ML/DL models, network, and node capabilities. The metrics included model accuracy, network usage (MB exchanged), and CPU usage (%), and the reported values represent the average across all nodes for each UC. As seen in Table 1, the Fedstellar platform showed a comparable accuracy rate of approximately 96% on average across all three use cases. The UC1 achieved a 90% accuracy rate in just 42 seconds, while the UC2 presented a significant improvement in network efficiency, using only 163 MB compared to 181 MB used by the UC1 (see Figure 3).

4 Conclusion

This demo paper presents Fedstellar, a platform for training FL models in a decentralized fashion. It offers a range of features and customization options, including the generation of federation topologies, architectures, ML/DL customization, simulated and physical deployment, and real-time monitoring. The platform demonstrated its capability and effectiveness through three use cases. However, Fedstellar is still under development, and there are many possibilities for further improvements. In this sense, the platform could be expanded to include more advanced FL techniques such as Vertical Federated Learning (VFL) or Trusted FL, allowing for even more secure and reliable training of models. In future research, a comprehensive analysis of advanced FL techniques and algorithms will be undertaken, assessing their performance and effectiveness across various application scenarios.

Acknowledgments

This work has been partially supported by *(a)* 21629/FPI/21, Fundación Séneca, *(b)* the Strategic Cybersecurity Project in Spain entitled CDL-TALENTUM (Development of Professionals and Researchers in Cybersecurity, Cyberdefence and Data Science) with the support of INCIBE and the European Mechanism for Recovery and Resilience (MRR) and as part of the measures of the Recovery, Transformation and Resilience Plan, *(c)* the Swiss Federal Office for Defense Procurement (armasuisse) with the DEFENDIS and CyberForce projects, and *(d)* the University of Zürich UZH.

References

[Barbieri *et al.*, 2022] Luca Barbieri, Stefano Savazzi, Mattia Brambilla, and Monica Nicoli. Decentralized federated learning for extended sensing in 6G connected vehicles. *Vehicular Communications*, 33:100396, 2022.

[Beltrán *et al.*, 2022] Enrique Tomás Martínez Beltrán, Mario Quiles Pérez, Pedro Miguel Sánchez Sánchez, Sergio López Bernal, Gérôme Bovet, Manuel Gil Pérez, Gregorio Martínez Pérez, and Alberto Huertas Celdrán. Decentralized federated learning: Fundamentals, state-of-the-art, frameworks, trends, and challenges. *arXiv preprint arXiv:2211.08413*, 2022.

[Beltrán *et al.*, 2023] Enrique Tomás Martínez Beltrán, Pedro Miguel Sánchez Sánchez, Sergio López Bernal, Gérôme Bovet, Manuel Gil Pérez, Gregorio Martínez Pérez, and Alberto Huertas Celdrán. Fedstellar Platform: Web Application, 2023. Available at http://federatedlearning.inf.um.es, Credentials - user: fedstellar; password: test123. Last accessed on 15/02/23.

[Georgatos *et al.*, 2022] Evangelos Georgatos, Christos Mavrokefalidis, and Kostas Berberidis. Efficient fully distributed federated learning with adaptive local links. *arXiv preprint arXiv:2203.12281*, 2022.

[Hard *et al.*, 2021] Andrew Hard, Kurt Partridge, Rajiv Mathews, and Sean Augenstein. Jointly learning from decentralized (federated) and centralized data to mitigate distribution shift. In *Proceedings of NeurIPS Workshop on Distribution Shifts*, 2021.

[Li *et al.*, 2020] Li Li, Yuxi Fan, Mike Tse, and Kuo-Yi Lin. A review of applications in federated learning. *Computers & Industrial Engineering*, 149:106854, 2020.

[McMahan *et al.*, 2016] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*, 2016.

[Nguyen *et al.*, 2022] T. V. Nguyen, M. A. Dakka, S. M. Diakiw, M. D. VerMilyea, M. Perugini, J. M. M. Hall, and D. Perugini. A novel decentralized federated learning approach to train on globally distributed, poor quality, and protected private medical data. *Scientific Reports*, 12:8888, 2022.