Research internship (Master 2)

Leveraging DevOps and reconfiguration for adaptive federated learning with heterogeneous clients

Hélène Coullon
IMT Atlantique, LS2N, Inria, France

Kandaraj Piamrat
Nantes Université, LS2N, Inria, France

2022-2023

Keywords *:* dynamic adaptation ; federated learning ; client selection.

1 Context

With the emergence of IoT devices, huge amount of data are generated continuously, which offers great opportunities for optimization, automation, and help in our day-by-day decisions (smart-everything, Industry 4.0, etc.). One way of taking advantage of this knowledge is to use machine learning (ML) techniques where data are used to train a model that is then able to make decisions according to the situation. However, at the same time, some privacy issues arise related to the transfer and leakage of personal, sometimes sensitive, information from end devices to central training entity in the Cloud. To handle this issue, Federated Learning (FL) has been proposed to provide a collaborative learning mechanism, which allows multiple parties to build a joint machine learning model. FL allows devices to keep private data locally where they were generated (within their controlled domains). Only local model updates are shared, typically, to a central aggregation server hosted by one of the parties or by a cloud service provider.

Classical FL architecture is composed of two types of distributed software components : the server ; and the associated clients. The FL server first communicates the initial global model to a set of selected FL clients. Clients perform a training with local data during a certain number of epochs in order to obtain the local models. Then, the clients send back their models to the central server for aggregation after a certain number of rounds. The mechanism of client selection and scheduling are particularly crucial in the learning process. Indeed, the server needs to select a subset of clients to distribute the global model. This impacts the learning performance. Clients selection is a difficult task that depends on different criteria such as system related computational, storage, and connection capacities, energy consumption, or statistical criteria such as the IID or Non-IID data, etc.

More generally, tuning configuration for optimizing an FL performance (e.g., time-to-accuracy, test accuracy) is a complex task. There are many ways of optimizing and parameterize an FL process, for example, changing the FL mechanism (horizontal, vertical, etc.) ; changing the way privacy is guaranteed (differential privacy, homomorphic encryption, etc.) ; selecting the appropriate model ; changing hyper-parameters of the models ; changing and selecting the set of clients involved in a training round, etc. **In this internship, we are interested more particularly in the client selection mechanism.**

As illustrated in the literature, client selection is difficult because of clients' heterogeneity (both in terms of system and data). A recent work called PyramidFL tackles this problem in [4]. Some approaches go further and conduct a real experiment such as in [6]. In fact, the authors of FedMSA first proposes a model selection algorithm, through the selection of training tasks. Once selected, the set of tasks are automatically and efficiently deployed as microservices in the Cloud and the Edge computing.

While a manual or ad-hoc way of dynamically adapting an FL approach is one possible approach, another approach is to leverage the genericity of microservices, DevOps and reconfiguration techniques, well adapted to dynamic changes and management, to automatically, efficiently and safely handle those dynamic changes to any FL approach. This is the goal of FL platforms such as for instance KubeFATE [1], Akraino [2], and OpenFL [3]. However, **existing platforms are quite limited regarding dynamic adaptations and reconfigurations of FL approaches**, and are mainly restricted to the deployment automation.

2 Work

Our goal in this internship is to **further study how to leverage advanced DevOps and reconfiguration techniques to offer more advanced dynamic adaptation when using FL, in particular when optimizing heterogeneous clients selection.**

To this purpose, we first plan to study a specific case-study of FL that is proposed in our research team to handle attack detection [1] in industrial IoT (IIoT). In fact, in the context of IIoT environment, most of the time, various types of labeled and unlabeled data are available. In order to take advantage of both types of data without the burden of labeling all, a semi-supervised approach based on autoencoder (AE) called FLUIDS has been proposed. First, an AE is trained on each device (using unlabeled local/private data) to learn the representative and low-dimensional features. Then, a cloud server aggregates these models into a global AE using FL.

Second, in our research team, previous work has already been done on efficient and safe reconfiguration languages [3, 5, 2] which could be leveraged in the context of this internship. More generally, our research team is particularly interested by advanced DevOps techniques for microservices orchestration (Kubernetes), as well as *service mesh* solutions which offer even more advanced dynamic features (Istio) to reconfigure a distributed containerized system.

The internship will be structured as followed :

1. State of the art on existing FL platforms and academic platforms, how they handle the dynamic adaptation when using FL ;
2. Study and understand our FL case-study, and where and when a client selection is required for optimization purpose ;
3. Considering the reconfiguration case-study identified in the previous bullet, study how to leverage DevOps approaches (container orchestration, service mesh, etc.) or reconfiguration languages [3] to perform this reconfiguration ;
4. Study decision algorithms to choose the new set of clients according to some optimization criteria ;
5. Conduct experiments to validate the approach and compare to existing approaches [4] ;
6. Write a scientific paper or a scientific report of the results.

3 Expected skills

The following skills are expected from the successful candidate :

- a student in the last year of a Master's degree in Computer Science (or in the last year of an engineering school with a computer science option) ;
- knowledge and experience on distributed software systems, in particular micro-services ;

1. `https://github.com/FederatedAI/KubeFATE`
2. `https://wiki.akraino.org/display/AK/Federated+Learning`
3. `https://openfl.readthedocs.io/en/latest/source/openfl/communication.html`

- knowledge and experience on DevOps approaches such as Infrastructure-as-Code, containers, orchestration etc. ;
- knowledge and experience on Machine Learning, and ideally federated learning ;
- knowledge of the Python programming language ;
- a good level of English to contribute to the writing of a research paper ;
- an ability to collaborate and communicate ;
- curiosity and an appetite for learning new things.

4 Additional information

Advisors

— Hélène Coullon, IMT Atlantique & Inria team STACK, helene.coullon@imt-atlantique.fr

— Kandaraj Piamrat, Nantes Université & Inria team STACK, kandaraj.piamrat@univ-nantes.fr

Duration 6 months

Salary legal amount of 3,90€ / hour, full time

Location IMT Atlantique, équipe Inria Stack, laboratoire LS2N à Nantes

Références

[1] Ons Aouedi, Kandaraj Piamrat, Guillaume Muller, and Kamal Singh. Federated semi-supervised learning for attack detection in industrial internet of things. *IEEE Transactions on Industrial Informatics*, pages 1–1, 2022. doi:10.1109/TII.2022.3156642.

[2] Maverick Chardet, Hélène Coullon, and Christian Pérez. Predictable Efficiency for Reconfiguration of Service-Oriented Systems with Concerto. In *CCGrid 2020 : 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*, Melbourne, Australia, May 2020. IEEE. doi:10.1109/CCGrid49817.2020.00-59. URL `https://hal.inria.fr/hal-02535077`.

[3] Maverick Chardet, Hélène Coullon, and Simon Robillard. Toward Safe and Efficient Reconfiguration with Concerto. *Science of Computer Programming*, 203 :1–31, March 2021. doi:10.1016/j.scico.2020.102582. URL `https://hal.inria.fr/hal-03103714`.

[4] Chenning Li, Xiao Zeng, Mi Zhang, and Zhichao Cao. Pyramidfl : A fine-grained client selection framework for efficient federated learning. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, MobiCom '22, page 158–171, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391818. doi:10.1145/3495243.3517017. URL `https://doi.org/10.1145/3495243.3517017`.

[5] Simon Robillard and Hélène Coullon. SMT-Based Planning Synthesis for Distributed System Reconfigurations. In *FASE 2022 : 25th International Conference on Fundamental Approaches to Software Engineering*, Munich, Germany, April 2022. URL `https://hal.inria.fr/hal-03536643`.

[6] Rui Sun, Yinhao Li, Tejal Shah, Ringo W. H. Sham, Tomasz Szydlo, Bin Qian, Dhaval Thakker, and Rajiv Ranjan. Fedmsa : A model selection and adaptation system for federated learning. *Sensors*, 22(19), 2022. ISSN 1424-8220. doi:10.3390/s22197244. URL `https://www.mdpi.com/1424-8220/22/19/7244`.