

A Quick Introduction to Federated Learning

Methods, Challenges, and Applications

Ethan Young

UCLA

Stats 147: Data Technologies for Data Scientists
Winter 2023

What is Federated Learning?

Federated learning (FL) is a machine learning (ML) technique that trains an algorithm across multiple servers (nodes) holding local data, without exchanging them

What is Federated Learning?

There are 3 main FL settings¹:

- **Centralized**

- Central server coordinates the participating nodes in the learning process

- **Decentralized**

- Nodes coordinate themselves to train the global model


- **Heterogeneous²**

- Involves a large set of heterogeneous clients such as mobile phones and internet of things (IoT) devices

¹Kairouz et al., *Advances and Open Problems in Federated Learning*, 2019

²Diao, Ding, and Tarokh, *HeteroFL: Computation and Communication Efficient Federated Learning for Heterogeneous Clients*, 2020

What is Federated Learning?



Why Federated Learning?

The appeal of FL lies in building a robust ML model without sharing data, which helps to address major societal concerns such as data privacy, data security, and access to data³

³Li et al., “A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection”, 2021

Learning Procedure

A summary⁴ of the learning procedure in the centralized setting is as follows:

- 1 Initialization
- 2 Client selection
- 3 Configuration
- 4 Reporting
- 5 Termination

⁴Bonawitz et al., *Towards Federated Learning at Scale: System Design*, 2019

Initialization

1 Initialization

- Initialize a ML model to be trained on the nodes

2 Client selection

3 Configuration

4 Reporting

5 Termination

Client Selection

- 1 Initialization
- 2 **Client selection**
 - A fraction of nodes are selected to start training on local data
- 3 Configuration
- 4 Reporting
- 5 Termination

Configuration

- 1 Initialization
- 2 Client selection
- 3 **Configuration**
 - Central server orders selected nodes to train the model on their local data
- 4 Reporting
- 5 Termination

Reporting

- 1 Initialization
- 2 Client selection
- 3 Configuration
- 4 **Reporting**
 - Selected nodes send their models to the central server for aggregation
 - Central server aggregates the models and sends updated model to nodes
 - Next round starts by returning to client selection
- 5 Termination

Termination

- 1 Initialization
- 2 Client selection
- 3 Configuration
- 4 Reporting
- 5 **Termination**
 - When a pre-specified criterion is met the central server aggregates the updated models into the final global model

Popular Variations

- There are an overwhelming number of FL variations that exist to address issues such as heterogeneous data and non-IID data
- Below are 2 that are widely-used as benchmarks and serve as the foundation of many other FL frameworks include:
 - **Federated stochastic gradient descent**⁵ (FedSGD)
 - **Federated averaging**⁶ (FedAvg)

⁵Shokri and Shmatikov, “Privacy-Preserving Deep Learning”, 2015

⁶McMahan et al., *Communication-Efficient Learning of Deep Networks from Decentralized Data*, 2016

Recent Advances

- It is difficult to know which methods are preferable to others due to constant developments in the field
- Below are several (of many) relatively recent advances that have captured much attention:
 - **Inverse Distance Aggregation**⁷ (IDA)
 - **FL with Dynamic Regularization**⁸ (FedDyn)
 - **Hybrid Federated Dual Coordinate Ascent**⁹ (HyFDCA)

⁷Yeganeh et al., *Inverse Distance Aggregation for Federated Learning with Non-IID Data*, 2020

⁸Acar et al., *Federated Learning Based on Dynamic Regularization*, 2021

⁹Overman, Blum, and Klabjan, *A Primal-Dual Algorithm for Hybrid Federated Learning*, 2022

Adversarial Attacks

- Understanding the impact of malicious actors (attackers) is a major challenge to the robustness of models learned by FL
- Chen et al.¹⁰ details several types of attacks on different aspects of FL
 - Examples include **Byzantine attacks**, **reconstruction attacks**, and **poisoning attacks**

¹⁰Chen et al., *Federated Learning Attacks and Defenses: A Survey*, 2022

Defense Mechanisms

- An active area of FL research is developing defense methods to prevent data breaches
- Chen et al.¹¹ describes 2 different levels of defenses:
security-based and **privacy-based**
 - Examples include **data anonymization**, **differential privacy**, and **secure multi-party computation**

¹¹Chen et al., *Federated Learning Attacks and Defenses: A Survey*, 2022

Healthcare

- The capacity for FL to address challenges of data privacy makes it a powerful tool for ML applications in healthcare
- **Partial meta-federated learning**¹² (PMFL) shows great potential in its fast training speed and high accuracy when applied to heterogeneous medical records

¹²Zhang et al., *PMFL: Partial Meta-Federated Learning for heterogeneous tasks and its applications on real-world medical records*, 2021

Satellite Constellations

- Low Earth Orbit (LEO) constellations that contain many satellites have become a large data source, although that data is expensive and slow to transfer
- **Asynchronous federated learning for LEO satellite constellations**¹³ (AsyncFLEO) outperforms existing methods by increasing convergence time and model accuracy

¹³Elmahallawy and Luo, *AsyncFLEO: Asynchronous Federated Learning for LEO Satellite Constellations with High-Altitude Platforms*, 2022

Internet of Things

- ML models are increasingly popular in industrial settings due to sensor data from production machinery becoming more widely available
- **Autoencoder-based federated learning**¹⁴ applied to sensor data reduced the network usage and demonstrates the success of FL in the industrial IoT

¹⁴Becker et al., *Federated Learning for Autoencoder-based Condition Monitoring in the Industrial Internet of Things*, 2022

Additional Resources

- **Federated Learning**¹⁵ is an in-depth exploration of relevant challenges and methods in FL
- **FedML**¹⁶ is an open-source and collaborative research library that supports FL algorithm development
 - Other libraries mentioned by the authors include **TensorFlow Federated** (TFF) and **PySyft**
- **OpenFL**¹⁷ is another open-source framework with TensorFlow and PyTorch training pipelines

¹⁵Ludwig and Baracaldo, *Federated Learning*, 2022

¹⁶He et al., *FedML: A Research Library and Benchmark for Federated Machine Learning*, 2020

¹⁷Foley et al., *OpenFL: the open federated learning library*, 2022

Current Landscape

- While FL circumvents a number of issues faced by traditional, centralized ML approaches, many open problems remain¹⁸
- These include node trustworthiness, robustness to adversarial attacks, improvements to communication efficiency, and development of privacy-preserving techniques

¹⁸Kairouz et al., *Advances and Open Problems in Federated Learning*, 2019

References I

-  Acar, Durmus Alp Emre et al. *Federated Learning Based on Dynamic Regularization*. 2021. DOI: [10.48550/ARXIV.2111.04263](https://doi.org/10.48550/ARXIV.2111.04263).
-  Becker, Soeren et al. *Federated Learning for Autoencoder-based Condition Monitoring in the Industrial Internet of Things*. 2022. DOI: [10.48550/ARXIV.2211.07619](https://doi.org/10.48550/ARXIV.2211.07619).
-  Bonawitz, Keith et al. *Towards Federated Learning at Scale: System Design*. 2019. DOI: [10.48550/ARXIV.1902.01046](https://doi.org/10.48550/ARXIV.1902.01046).
-  Chen, Yao et al. *Federated Learning Attacks and Defenses: A Survey*. 2022. DOI: [10.48550/ARXIV.2211.14952](https://doi.org/10.48550/ARXIV.2211.14952).
-  Diao, Enmao, Jie Ding, and Vahid Tarokh. *HeteroFL: Computation and Communication Efficient Federated Learning for Heterogeneous Clients*. 2020. DOI: [10.48550/ARXIV.2010.01264](https://doi.org/10.48550/ARXIV.2010.01264).

References II

-  Elmahallawy, Mohamed and Tie Luo. *AsyncFLEO: Asynchronous Federated Learning for LEO Satellite Constellations with High-Altitude Platforms*. 2022. DOI: [10.48550/ARXIV.2212.11522](https://doi.org/10.48550/ARXIV.2212.11522).
-  Foley, Patrick et al. *OpenFL: the open federated learning library*. 2022. DOI: [10.1088/1361-6560/ac97d9](https://doi.org/10.1088/1361-6560/ac97d9).
-  He, Chaoyang et al. *FedML: A Research Library and Benchmark for Federated Machine Learning*. 2020. DOI: [10.48550/ARXIV.2007.13518](https://doi.org/10.48550/ARXIV.2007.13518).
-  Kairouz, Peter et al. *Advances and Open Problems in Federated Learning*. 2019. DOI: [10.48550/ARXIV.1912.04977](https://doi.org/10.48550/ARXIV.1912.04977).
-  Li, Qinbin et al. "A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection". In: *IEEE Transactions on Knowledge and Data Engineering* (2021). DOI: [10.1109/tkde.2021.3124599](https://doi.org/10.1109/tkde.2021.3124599).

References III



Ludwig, Heiko and Nathalie Baracaldo, eds. *Federated Learning. A Comprehensive Overview of Methods and Applications*. Springer Cham, 2022. DOI: <https://doi.org/10.1007/978-3-030-96896-0>.



McMahan, H. Brendan et al. *Communication-Efficient Learning of Deep Networks from Decentralized Data*. 2016. DOI: [10.48550/ARXIV.1602.05629](https://doi.org/10.48550/ARXIV.1602.05629).



Overman, Tom, Garrett Blum, and Diego Klabjan. A *Primal-Dual Algorithm for Hybrid Federated Learning*. 2022. DOI: [10.48550/ARXIV.2210.08106](https://doi.org/10.48550/ARXIV.2210.08106).



Shokri, Reza and Vitaly Shmatikov. “Privacy-Preserving Deep Learning”. In: *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 2015. DOI: [10.1145/2810103.2813687](https://doi.org/10.1145/2810103.2813687).

References IV



Yeganeh, Yousef et al. *Inverse Distance Aggregation for Federated Learning with Non-IID Data*. 2020. DOI: [10.48550/ARXIV.2008.07665](https://doi.org/10.48550/ARXIV.2008.07665).



Zhang, Tianyi et al. *PMFL: Partial Meta-Federated Learning for heterogeneous tasks and its applications on real-world medical records*. 2021. DOI: [10.48550/ARXIV.2112.05321](https://doi.org/10.48550/ARXIV.2112.05321).