ICANN24

33rd International Conference on Artificial Neural Networks

UPPSALA UNIVERSITET

scaleout

SciML UU

Dalle Molle Institute for Artificial Intelligence

IDSIA

SUPSI

ENNS

AiChemist

AIDD

ICANN24 – Tutorial Session

FEDn - A scalable federated machine learning framework for cross-device and cross-silo environments

Introduction

Federated machine learning has opened new avenues for privacy-preserving data analysis. Instead of pooling data in a central location, different data owners or IoT devices keep data local and training is decentralized where only model parameters are exchanged. It is an active area of research where most of the current efforts focus on the algorithmic details and communication overhead required to train accurate models. Despite much progress in the field, production-grade federated machine learning frameworks that deal with fundamental properties such as scalability, fault tolerance, security and performance in geographically distributed settings have not been available to the ML-engineer. To fill this gap, Scaleout Systems and SciML at Uppsala University have designed and developed the *FEDn* framework. FEDn is an open-source framework dedicated to address federated machine learning challenges at scale. This tutorial aims to provide a knowledge sharing platform and highlight challenges and possible solutions related to federated machine learning. The tutorial will also cover hands-on experience with the *FEDn* framework, present the latest experiment results based on large numbers of clients running and jointly training models in a heterogeneous distributed environment.

Session type and duration

Tutorial with a hands-on session, 2 hours

Intended audience (introductory, intermediate, advanced)

Introductory and intermediate

Presenters

- Associate Professor Salman Toor
  - Email: Salman.Toor@it.uu.se, Uppsala University

Salman Toor is Associate Professor in Scientific Computing at Uppsala University. He is an expert in the field of distributed computing infrastructures and applied machine learning. Toor is the co-chair of the Scientific Machine learning (SciML) research group at Uppsala University and co-founder and CTO at Scaleout Systems AB. From the platform of Scaleout Systems, Toor has organized a number of national and international workshops and special sessions.

- Associate Professor Andreas Hellander
  - Email: Andreas.Hellander@it.uu.se, Uppsala University

Andreas Hellander is Associate Professor in computational science and engineering, co-chair of the Scientific Machine learning (SciML) research group at Uppsala University and CEO of Scaleout Systems. His scientific interests include federated machine learning and scalable digital experiments for complex systems, with 50+ scientific publications in scientific computing and its applications.

- Responsible Research Groups –
  - Scientific Machine Learning (SciML) https://sciml.se/
  - Scaleout Systems (Scaleout) https://www.scaleoutsystems.com/

Agenda of the special session/tutorial

Total Duration 2 hours

1. Introduction (20 minutes)
    1. Introduction to federated machine learning
    2. Challenges related to Federated machine learning
    3. Different architectures (central, hierarchical and fully distributed)
    4. Well-known frameworks of federated machine learning
2. FEDn Framework (20 minutes)
    1. Design philosophy and architecture
    2. Implementation details
    3. Results based on cross-silo and cross-device use cases
3. Discussion session (20 minutes)
4. Hands-on Session (50 minutes)
5. Summary and closing remarks (5 minutes)

Requirements for the hands-on session

- Attendees need to have a stable internet connection.
- Hardware requirements:
    - Memory requirement, more than 4GB
    - Standard four core physical or virtual machine
    - Storage requirement, 5GB
- Software requirement:
    - Access to a Linux/windows/Mac machine, physical or virtual environment
    - A community version of Docker and Docker-compose environments
    - Python 3.10 or above
    - Preferably the latest Chrome browser

Prerequisite knowledge or skills required for attendees

- Introductory level understanding of neural networks
- Software
    - Basic understanding of the Linux command-line environment
    - Basic understanding of Docker containers
    - Intermediate-level Python programming skills
- Hardware
    - A laptop with a Linux/Windows/Mac and Docker environment

Base Article

Morgan Ekmefjord, Addi Ait-Mlouk, Sadi Alawadi, Mattias Åkesson, Prashant Singh, Ola Spjuth, Salman Toor, Andreas Hellander. *Scalable federated learning with FEDn*, to appear in the 2022 IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGrid). ArXiv preprint: https://arxiv.org/abs/2103.00148