ACM MobiCom 2021 Tutorial

Federated Mobile Sensing for Activity Recognition

Organisers:

Stefanos Laskaridis, Samsung AI Center Cambridge
Dimitris Spathis, University of Cambridge
Mario Almeida, rain.us

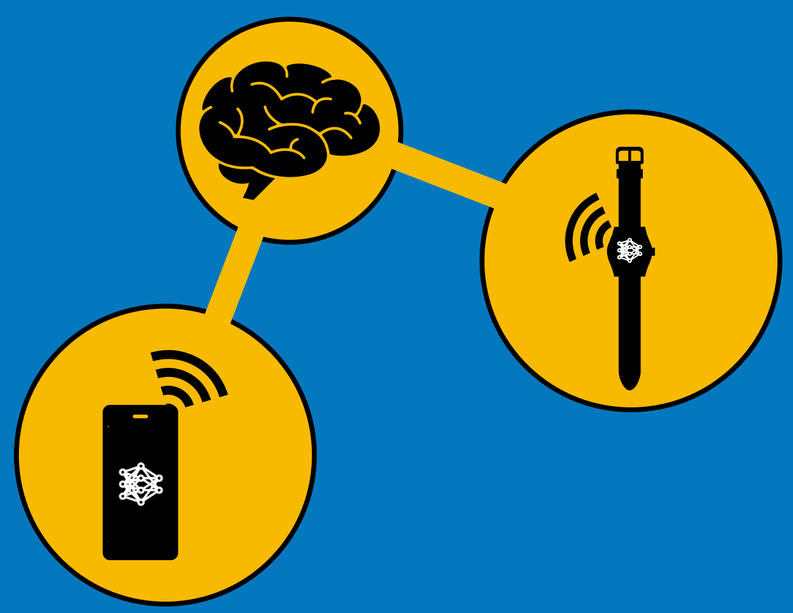

Abstract

Despite advances in hardware and software enabling faster on-device inference, training Deep Neural Networks (DNN) models has largely been a long-running task over TBs of collected user data in centralised repositories.

Federated Learning has emerged as an alternative, privacy-preserving paradigm to train models without accessing directly on-device data, by leveraging device resources to create per client updates and aggregate centrally. This has been applied to various tasks, ranging from next-word prediction to automatic speech recognition (ASR).

In this tutorial, we recognise on-device sensing as a privacy-sensitive task and build a federated learning system from scratch to showcase how to train a model for accelerometer-based activity recognition in a federated manner. In addition, we present the current landscape and challenges in the realm of federated learning and mobile sensing and provide guidelines on how to build such systems in a privacy-preserving and scalable manner.

Schedule

- Keynote #1: Federated Learning: Current landscape and challenges
- Keynote #2: On-device sensing
- Tutorial: Federated Sensing hands-on
- Keynote #3: Scaling federated training on distributed infrastructures

The tutorial is co-located with ACM MobiCom'21, in the beautiful city of New Orleans. Hybrid participation is available.

Conference Dates: March 28th to April 1st, 2022

Tutorial Date: April 1st, 2022

https://federatedsensing.gitlab.io