Internship project description
Feerated learning with untrusted server

Jan Ramon

10/11/2022

1 Motivation and context

Over the last decades, there has been an increasing interest in exploiting data. On the other hand, recently there has also been an increasing awareness of the risks of collecting sensitive data centrally, given the frequency of data leaks, hacking or abuse. INRIA's Magnet team is interested in decentralized privacy-preserving machine learning where the sensitive data remains with the data owners, and machine learning is performed collaboratively by these data owners by participating in collaborative algorithms which (through the use of differential privacy and/or encryption) generate the desired statistical models but prevents sensitive data from being revealed. Important ongoing research projects in the team include the TIP, TRUMPET and FLUTE projects. This internship fits into the larger research program including these projects.

Most current proposals for federated learning involve one or more servers which either have a trusted execution environment (TEE) or are assumed to not collude (exchange information). This isn't fully satisfactory. First, in the current geopolitical environment governments are unwilling to trust TEEs on processors constructed by large companies based in other continents. Second, as over recent years media has reported about a large number of data leaks and about the huge economic interests in obtaining information, it is hard for consumers to trust that a set of servers don't collude, especially as such assumption is very difficult to verify.

2 Objectives

Therefore, the goal of this internship project is to build a Federated Learning system with Untrusted Servers (FLUS), in particular for scenarios where a server is not trusted to properly protect secrets.

Simplifying assumptions. Of course, a purely decentralized approach is sometimes very inefficient, hence we will assume that a data user, e.g., a company who has interest in the success of a machine learning effort, may have

an incentive to set up a server which properly coordinates a computation and organizes communication between parties. We will also assume that there is a proper public key infrastructure (PKI) so parties can prove they are unique citizens.

Prior work. In the past, the MAGNET team has researched this scenario and developed promising algorithms for it, e.g., in the context of the PhD thesis of Cesar Sabater. Also, in the TAILED project the MAGNET team is developing an open source library containing infrastructure to support developing secure, privacy-preserving machine learning algorithms. The road is now free to develop our first FLUS algorithms and experimentally evaluate them.

Objectives. In particular, the goal of the internship project is

- to develop a FLUS algorithm (probably for logistic regression, but we may change to another machine learning task if our medical partners express strong preferences before the start of the internship project)
- to experimentally evaluate its correctness and scalability on a medical benchmark dataset
- optionally, if time allows, to develop a more programmer-friendly interface to the system so developers can run scikit-learn style scripts as FLUS algorithms on decentralized data.

Whether work on the third optional objective is started depends on the skills of the student and the challenges encountered underway. Both with and without this part an interesting project can be carried out.

3 Plan

Here is a tentative work plan:

- Literature study. Among others, getting familiar with distributed algorithms, basic multi-party computing concepts, the existing TAILED library, related work on federated machine learning, logistic regression (2 weeks)
- Initial development cycle
  - Implementation of basic statistics aggregation strategies (6 weeks)
  - Implementation of a FLUS logistic regression algorithm (3 weeks)
  - Testing correctness and scalability on benchmark dataset (3 weeks)
- Extending features to improve privacy guarantees and scalability and to add features requested by users/partners. (3 weeks)

- Development of programmer-friendly interface, 'compiling' scikit-learn scripts into FLUS operations (3 weeks)
- Testing (4 weeks)
- Completion of the internship report (2 weeks)

4 Environment

The project will be conducted in the INRIA MAGNET team. The student will collaborate and interact with various other collaborators in the team, who are working on other but related component of our software. It is possible the student will be asked to present progress to the team or to partners in ongoing projects. An application for ZRR access will need to be made to the FSD.

5 Requirements

We expect the student has a good understanding of basic computer science concepts, e.g., data structures and algorithms, and at least some initial notions of distributed systems, cryptography and/or statistics. Ideally the project will involve a significant part in C++ and a smaller part in Python, but it can be adapted for students knowing only one of both languages.