Privacy Analysis of Federated Learning with Gradient Boosting Decision Trees*

Saloni Kwatra[0000−0002−4896−7849] and Vicenç Torra[0000−0002−0368−8037]

Department of Computing Science, Umeå University, Sweden
{salonik,vtorra}@cs.umu.se

Abstract. The goal of Federated Learning (FL) is to provide a privacy-preserving framework that allows devices to jointly learn a shared model while retaining all training data on the device, avoiding the need to store data in a centralized database. We study Weighted Gradient Boosting Decision Trees (WGBDT) in an FL framework. The concept of WGBDT is that a record is important if it is similar to many records of other devices. Therefore, an aggregated gradient of all similar records is used for updating the gradients when creating sequential decisions for Gradient Boosting Decision Trees (GBDT) for each device. One of the methods for determining the similar records among records from different devices is the Locality Sensitive Hashing (LSH), which places similar records in the same bucket, and dissimilar records in different buckets so that FL-based WGBDT can use the aggregated gradient of similar samples during the training of a GBDT model for each distributed device. One of the disadvantages of adopting the LSH approach for the FL-based WGBDT is that it assumes that each device knows the hash values of other devices' records for finding similar records for the computation of aggregated gradient. We demonstrate two data reconstruction attacks, called Least Squares (LS) based attack and a Non-Linear Optimization (NLO) attack. The purpose of our attacks is to re-engineer the original data of users on knowing the hashed values. According to our findings, when the number of LSH functions is more than the data dimension, we can reconstruct more than 90% of the training data, and when the number of hash functions (L) is smaller than the data dimension (d), it is still possible to reconstruct the data significantly. We also did data reconstruction attacks when the data is anonymised using Mondrian k-anonymity before applying LSH to obtain similar samples. We discovered that when the data is anonymised before the hash values are computed, the reconstruction accuracy is lower.

Keywords: Federated Learning · Privacy · Gradient Boosting Decision Trees · Locality Sensitive Hashing · Reconstruction Attacks.

* This study was partially funded by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.