Game of Gradients: Mitigating Irrelevant Clients in Federated Learning

Lokesh Nagalapatti, Ramasuri Narayanam

IBM Research-India

Federated Learning (FL)

Learn a Machine Learning model at a central server with contributions from data possessed by multiple clients (without sharing data).

- The server C trains a classifier $f(\theta)$, where θ is obtained by distributed training and aggregation over K clients.
- The hope is that it generalizes on the test dataset D_{test}.

FL Objective

$$\min_{\theta} l(\theta) = \sum_{k=1}^{K} \frac{n_k}{n} l_k(\theta);$$

$$\text{where } l_k(\theta) = \frac{1}{n_k} \Sigma_{i \in \mathbb{D}_k} l_i(\theta) \quad (1)$$

Challenges in Federated Learning

- Unbalanced data,
- non-IID (Independent and Identically Distributed),
- Participating clients are connected to the server with limited communication bandwidth,
- Privacy of the clients' data.

Noisy Clients in FedAvg

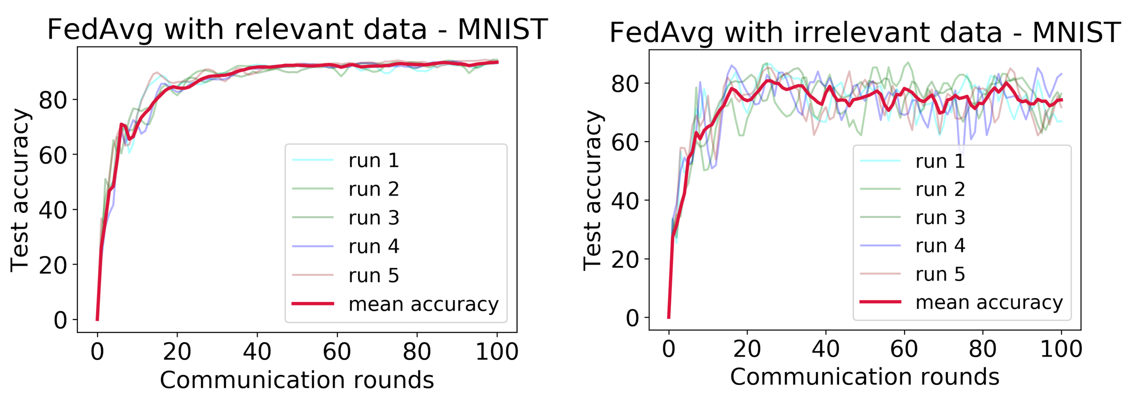

Problem Statement

- Develop a mechanism that facilitates the server to prune irrelevant clients.
- Server has no visibility into how each client generates the updates
- Server should evaluate the relevance of a client only working with the updates it receives from clients.
- We refer to these as Federated Relevant Client Selection (**FRCS**) problems.

Solution Approach

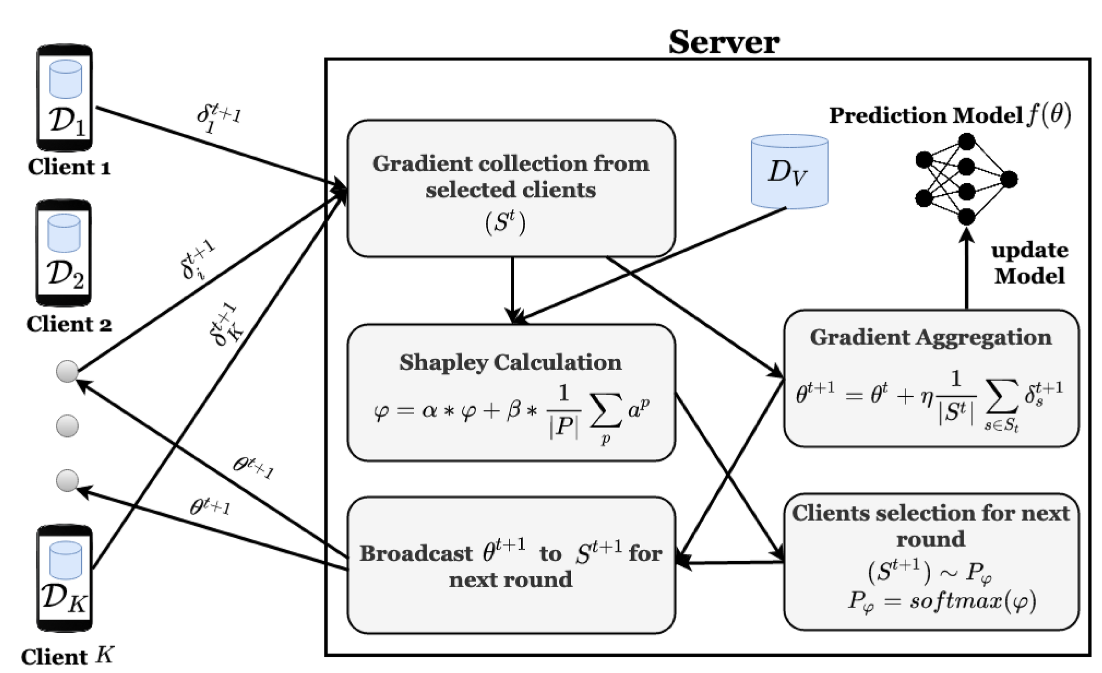

Figure 1:FRCS Architecture

Shapley values of Clients

Char. fn - $v(X) : 2^{\delta^{t+1}} \to \mathcal{R}$

$$\theta_X^{t+1} = \theta^t + \frac{1}{|X|} \sum_{s \in X} \delta_s^{t+1}$$

$$v(X, D_V) = \mathcal{P}(f_{\theta_X^{t+1}}, D_V)$$

$\mathcal{P}$ function denotes the performance of the central model with parameters θ_X^{t+1} on the validation data D_V

FRCS Problems

- Selecting relevant clients
- Selecting Class-label dependent relevant clients
- Detecting label corruption (upto permutation) and standardizing corrupted labels

FRCS: Relevant clients

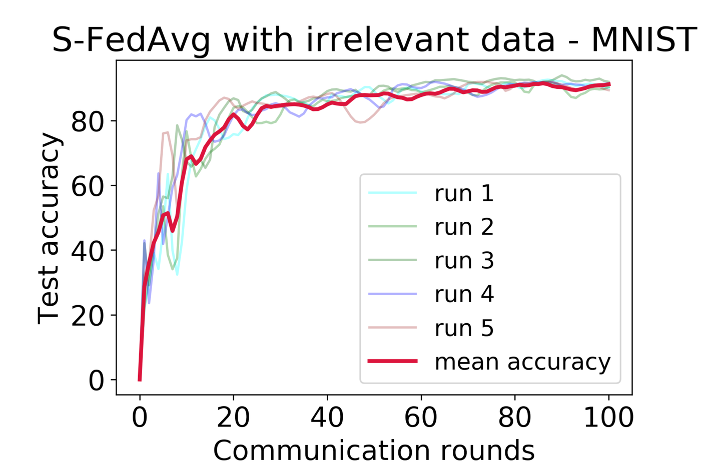

Figure 2:Accuracy

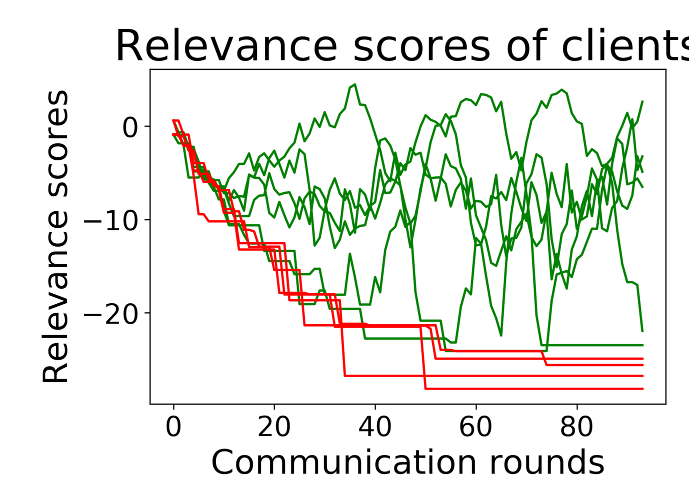

Figure 3:Relevance Scores (Shapley Values)

FRCS: Class dep. Client Selection

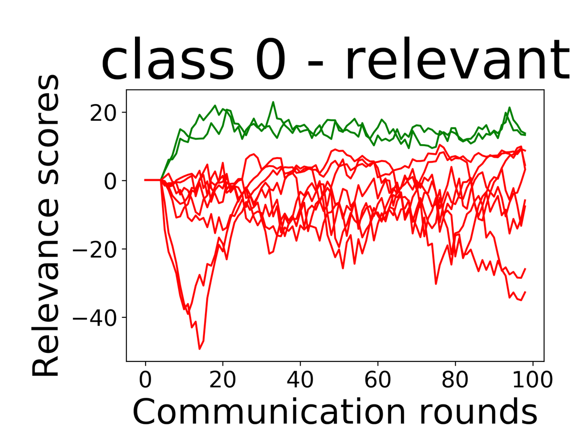

Class 0 - Relevant Data

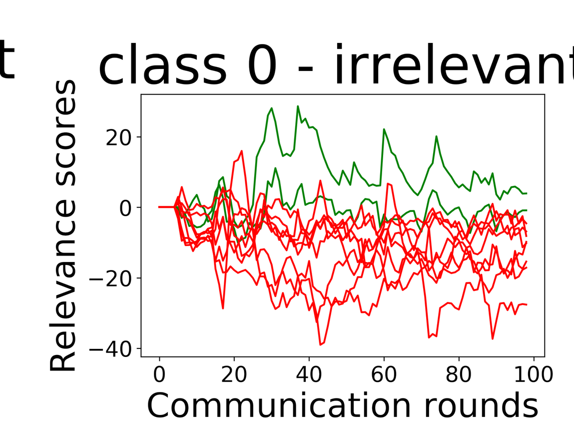

Class 0 - Irrelevant Data

FRCS: Data Label Standardization

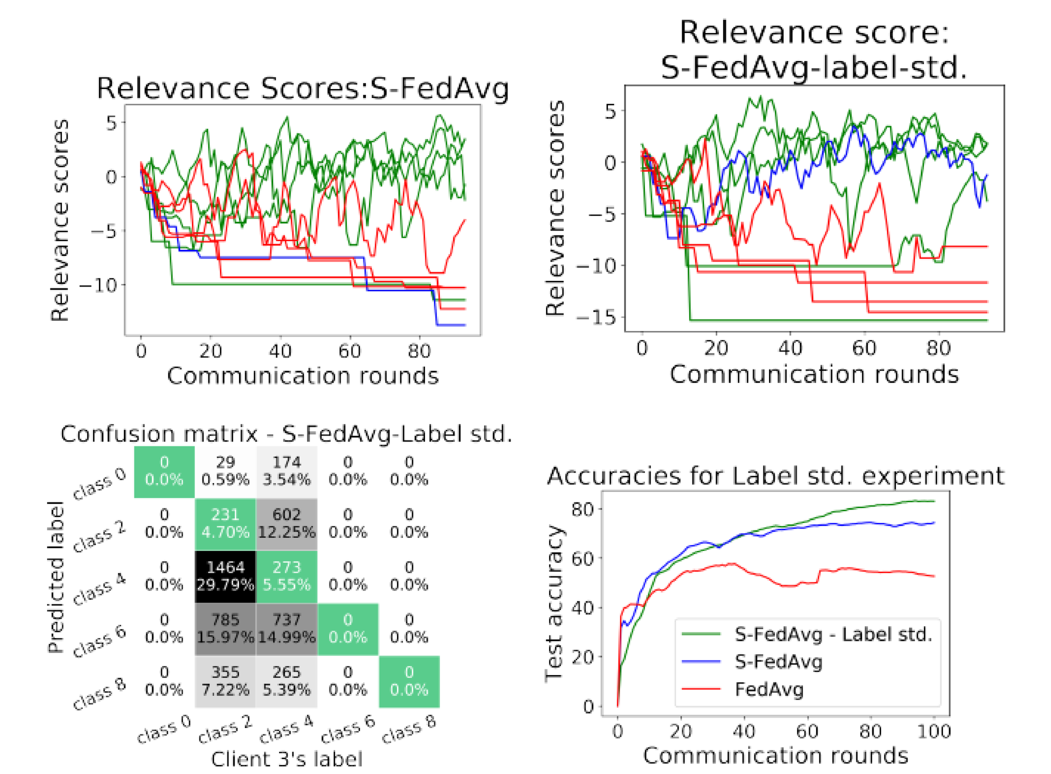

Figure 4:Error correction occurs at round 17 for the blue client

Conclusion

- Studied the impact of noisy clients in Federated Learning and proposed solution approaches to 3 important FRCS problems.
- At the heart of SFedAvg is Shapley value Computation
  - Scales exponentially with number of players
  - Hence we adopt Monte-Carlo approximation
- The approach is applicable as is in Cross-Silo Federated Learning (involving 10s of 100s of clients)

Contact Information

- Lokesh Nagalapatti; nlokeshiisc@gmail.com
- Ramasui Narayanam; ramasurn@in.ibm.com