Year 7 Proposal - Privacy Preserving Multi-Party Analytics

Stephen Adams (PI)[1], Steven Boker (PI)[1], Peter Beling (PI)[1], Steven Greenspan (IAB Project Mentor)[2]
University of Virginia[1], CA Technologies[2]

Project Start:	End Date:	Project Budget:	Spent:
08/31/2018	05/31/2019		

Project Summary:

Preserving privacy in machine learning on multi-party datasets is of importance to many domains. Existing solutions suffer from several critical limitations, such as significantly reduced utility after enforcing differential privacy, excessive communications burden between information fusion center and local participants who contribute data, etc, which severely limit their practical adoption. This project is a continuation of our year 6 project and consists of two parts. The first part is an extension of the differential private deep neural network model developed in project 6a.052.UVA under the multi-party heterogeneous dataset setting (see figure 1). In addition, we introduce the theoretical work in adaptive privacy budget into the setting. The second part shifts the focus to the paradigm of transfer learning, where the private datasets in the "sources" are different from the dataset in the "target" and the "target" aims to take advantage of the private knowledge transferred from sources to improve its own model (as shown in figure 2). More specifically, we are interested in the one-shot model aggregation technique and its application in the transfer learning scenario. This method will provide a practical solution to the multi-party privacy preserving deep learning under the transfer learning settings, which is especially beneficial to financial institutions who are willing to jointly learning machine learning models, but are prohibited by privacy restriction and lack sufficient labeled dataset. For example, if one "target" bank would like to train a fraud detection model but lacks labeled transactions, our method will ensure that the "target" bank can benefit from the models from other "source" banks based on how related their transactions are without violating their privacy.

Details of Progress/Achievements:
Literature review; dataset cleanup;

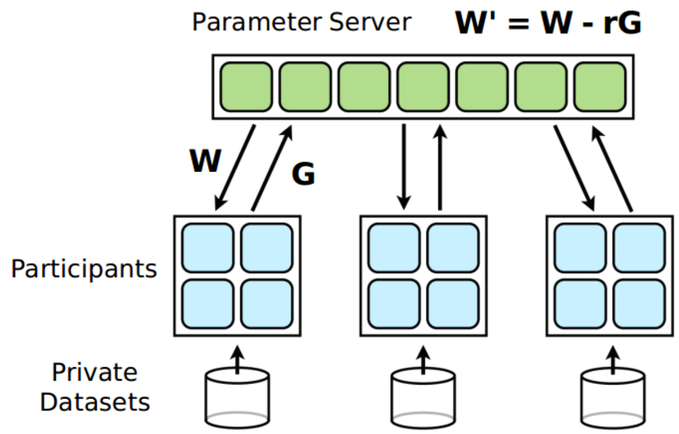

Figure 1. Iterative training process of the differential private deep neural network

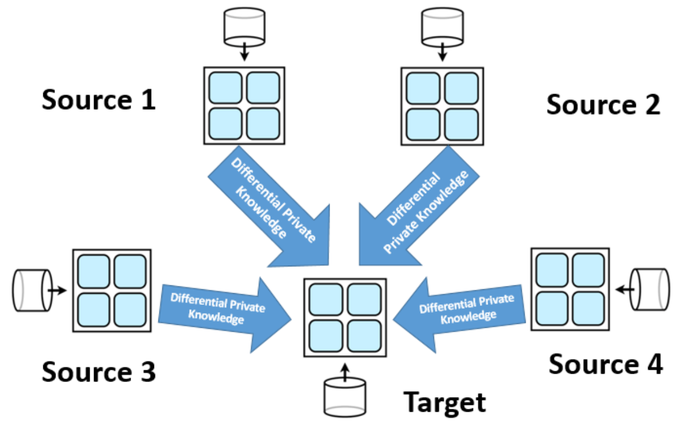

Figure 2. Model aggregation and transfer learning diagram

PROJECT DELIVERABLES

Deliverable	Achievements	Remaining To Do
1. Research paper on the algorithms and theories of distributed privacy preserving deep learning model.	Literature review	Algorithm design and utility guarantee.
2. Empirical evaluation on the utility-privacy trade-off of the method on real-world banking data and health-care data	Dataset cleanup	Feature engineering, method implementation