Private Cross-Silo Federated Learning For Extracting Vaccine Adverse Event Mentions

Pallika Kanani, Virendra J. Marathe
Daniel Peterson, Rave Harpaz,
Steve Bright

ORACLE Labs

ORACLE Health Sciences

Personalization helps recover the accuracy lost by enforcing differential privacy in Federated Learning

Automatically extracting mentions of suspected drug or vaccine adverse events (potential side effects) from unstructured text is critical in the current pandemic, but small amounts of labeled training data remains silo-ed across organizations due to privacy concerns. Federated Learning (FL) allows such users to jointly train a more accurate global model without physically sharing their data. We study this in the context of Named Entity Recognition (NER) task for a vaccine adverse event detection application. We ask the following questions as part of this study:

- Does *FL* perform better than *individual* models across users?
- Does accuracy drop when differential privacy (DP) is introduced?
- Does personalization help improve accuracy over *FL* and mitigate *DP-FL*’s accuracy loss enough to re-incentivize users to participate in the federation?

In the paper, we also show results on robustness to varying parameters of DP and stability against uncertainities of real world, such as users dropping out.

NER for the Vaccine Adverse Event Detection Reporting System (VAERS) Dataset

The prominent surveillance system for vaccines is the U.S. Centers for Disease Control and Prevention (CDC) and the Food and Drug Administration (FDA) Vaccine adverse Event reporting System (VAERS).

Each VAERS report includes a textual narrative describing the adverse event, along with other metadata, for example:

“Approximately 5 minutes after being given immunizations listed below pt started to cough and wheeze”

Named Entity Recognition (NER) Task: Automatically extract mentions of adverse events from unstructured data

We use a BiLSTM architecture, which consists of three major components: (1) a word representation layer made of word embeddings, (2) two stacked layers of bidirectional long short-term memory (LSTM) cells, and (3) a feedforward layer that performs the final BIO sequence labeling.

Private Federated Learning with Fine Tuning

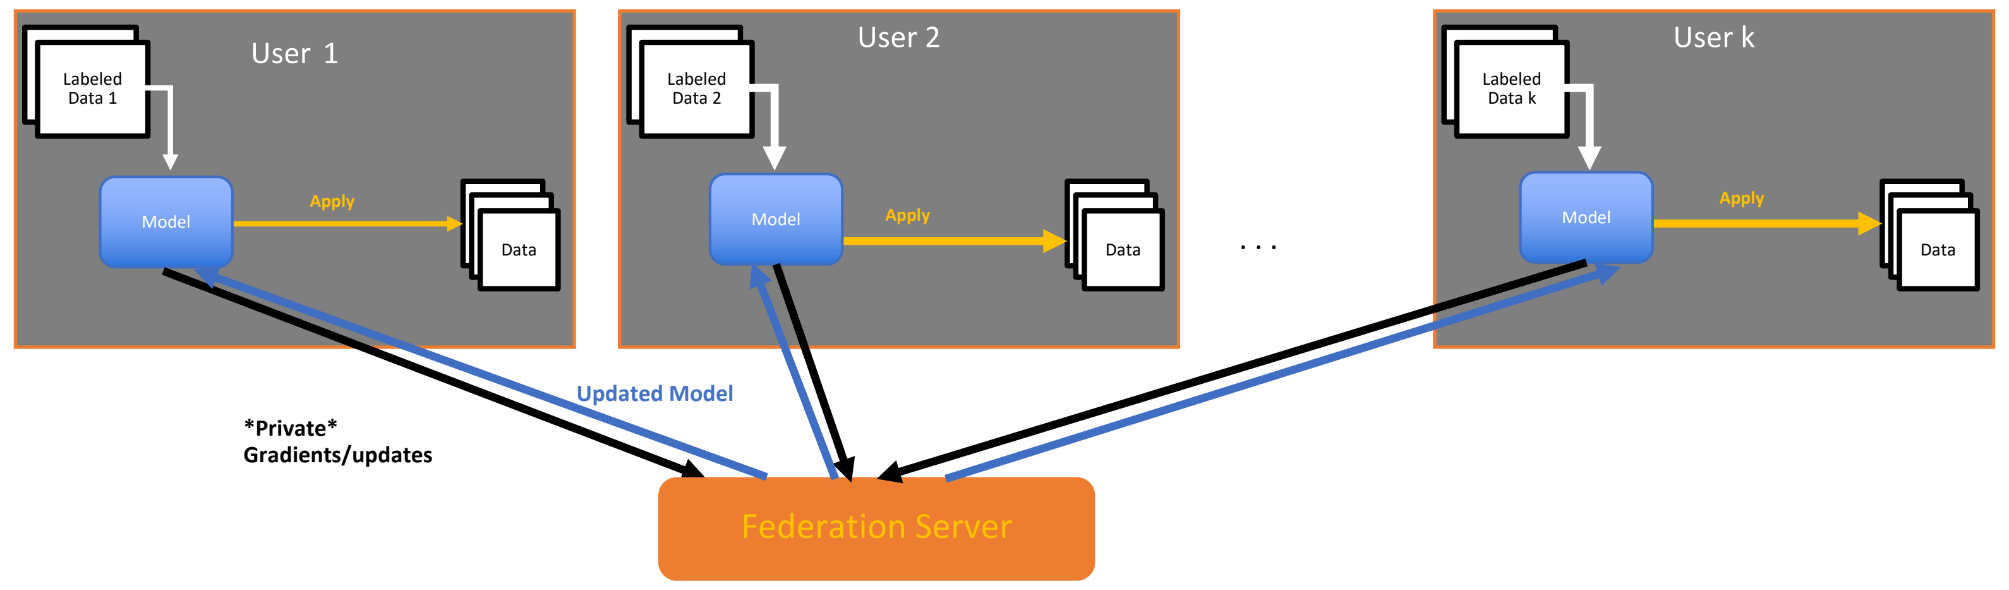

Architecture of Federated Learning System with Differential Privacy

For FL, we use one of the most widely used method of aggregation, FedAvg (McMahan et al. (2016)), where user parameters updates are averaged at the federation server and applied to the global model.

Noting privacy concerns, more recent work has proposed addition of differential privacy to FL. To enforce local DP, we use the algorithm proposed by Abadi et al. (2016) that injects gaussian noise (calculated using their moments accountant algorithm) in parameter gradients during local training at each user.

Researchers have recently proposed different forms of *personalization* approaches to remedy the problem of model degradation due to DP enforcement (Peterson et al. (2019); Yu et al. (2020)). We use FL with *Fine Tuning* Yu et al. (2020), and call the model FT-DP-FL.

Experimental Setup

The VAERS data (de-identified) is publicly available in structured format. For this study, we took ~18k reports from 2014-2017, and annotated it to create training data. We split these reports by the names of the manufacturers of the vaccine, and use 60% data for training, 10% for tuning, 10% for validation, and 20% for testing. We compute the precision, recall, and F1 of each token label on a 1-vs-all basis.

As the first baseline for our experiments, we train Individual models (*Ind*), i.e. assume that each manufacturer only uses their own training set, and test on their respective test set. This baseline represents the case in which the manufacturer chooses not to participate in the federation at all.

For FL, we use a learning rate of 0.01 and train all the federated models for 20 rounds of FedAvg, with additional 20 epochs for the fine tuning variants at each manufacturer. For enforcing DP, we use epsilon =2.0 and calculate the sigma values suitable per user.

Results

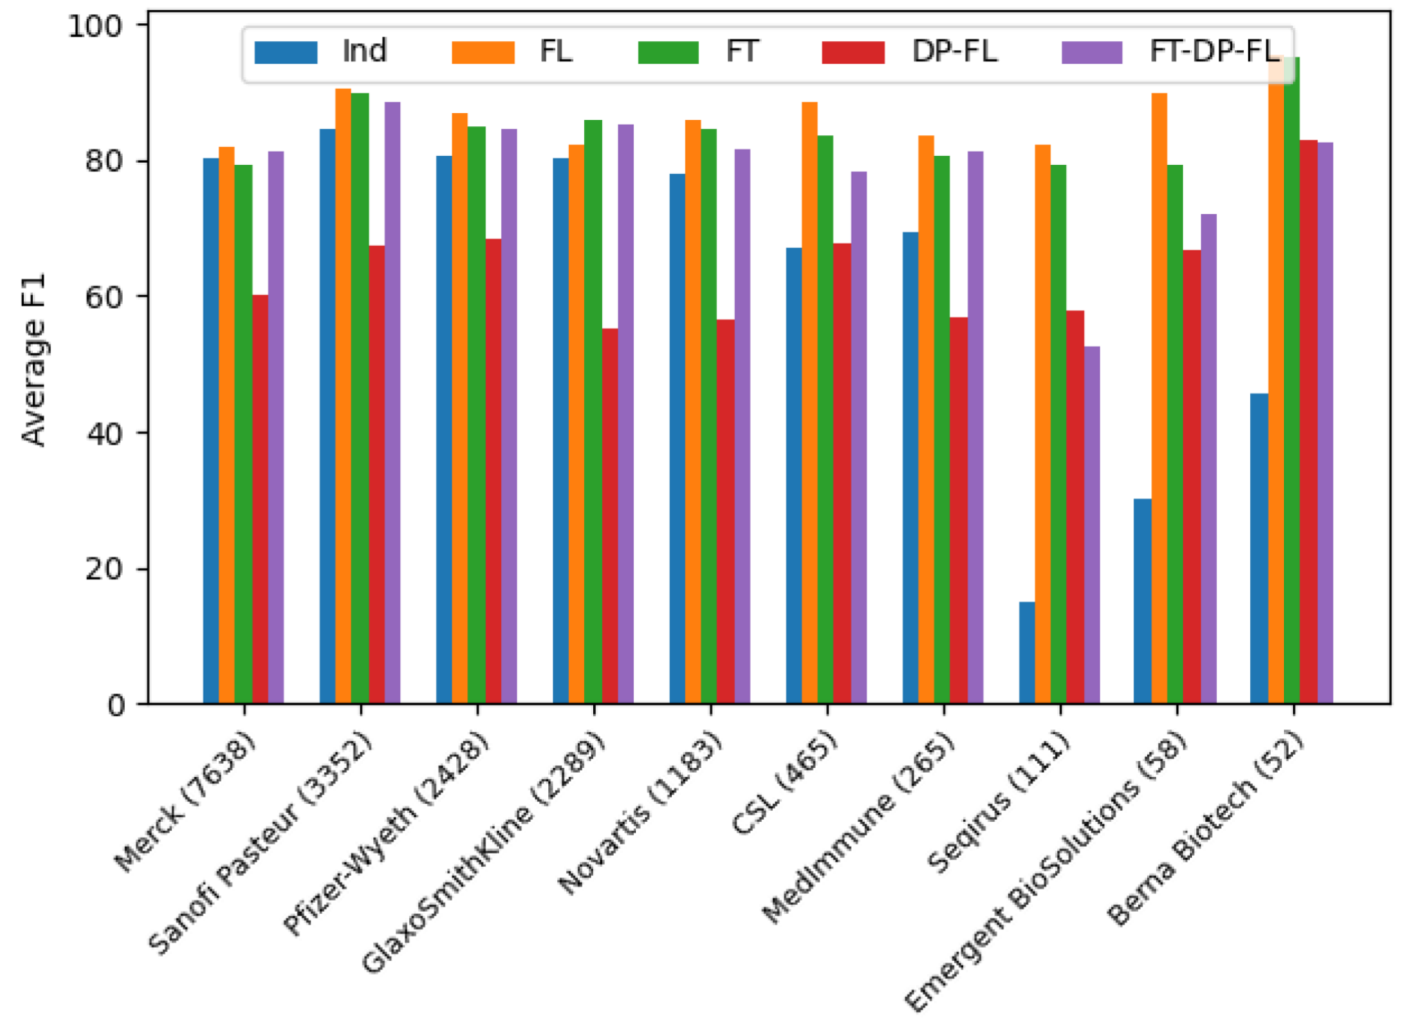

‘ Vaccine Manufacturer’ is a field in the public VAERS database that identifies the manufacturer of the vaccine reported in the VAERS form. There is no relationship between this field and the reporter. ‘Num VAERS Reports’ does not represent the rate of adverse events associated with the manufacturer or its products and cannot be used to estimate such rates. The statistics are based on a sample of reports submitted to VAERS between 2015-2017 whose MedDra coded adverse events appeared in the narrative. Because the statistics are based on a carefully selected sample, the distribution of reports shown may not represent the true distribution of reports associated with different vaccine manufacturers.

Vaccine Manufactuer	Individual F1	FL		FT-DP-FL	
		F1	Error Red.	F1	Error Red.
Merck Co. Inc.	80.10	82.00	9.55%	81.20	5.53%
Sanofi Pasteur	84.60	90.40	37.66%	88.40	24.68%
Pfizer-Wyeth	80.50	87.00	33.33%	84.60	21.03%
Glaxo-Smithkline Biologicals	80.20	82.20	10.10%	85.30	25.76%
Novartis Vaccines And Diagnostics	77.80	85.80	36.04%	81.50	16.67%
CSL Limited	67.10	88.50	65.05%	78.30	34.04%
Medimmune Vaccines Inc.	69.30	83.50	46.25%	81.10	38.44%
Seqirus Inc.	15.00	82.10	78.94%	52.60	44.24%
Emergent Biosolutions	30.10	89.70	85.26%	71.90	59.80%
Berna Biotech Ltd.	45.80	95.40	91.51%	82.50	67.71%

The FL model consistently outperforms *Individual* models for each of the users, including large manufacturers with a lot of training data. As we add noise related to differential privacy to the federated learning model, F1 values drop significantly across the board. Applying fine tuning in this case helps bring it back up to the point, where it is again advantageous for each party to participate in the federation.

Conclusion

Extracting mentions of vaccine adverse events using machine learning methods is an extremely urgent task right now. Federated Learning is a promising approach for breaking down organizational and geographical barriers to collaboration on building very effective models to solve this problem. Our work demonstrates that manufacturers with dataset of all different sizes can benefit from participating in such a federation, and that the loss of accuracy incurred through adding additional layers of privacy can be mitigated by introducing personalization.