

# Federated Deep Learning Architecture for Personalized Healthcare

Helen CHEN<sup>a,1</sup>, Shubhankar MOHAPATRA<sup>b</sup>, George MICHALOPOULOS<sup>b</sup>, Xi HE<sup>b</sup>  
and Ian MCKILLOP<sup>a,b</sup>

<sup>a</sup>*School of Public Health and Health Systems, University of Waterloo, Canada*

<sup>b</sup>*Cheriton School of Computer Science, University of Waterloo, Canada*

**Abstract.** Using deep learning to advance personalized healthcare requires data about patients to be collected and aggregated from disparate sources that often span institutions and geographies. Researchers regularly come face-to-face with legitimate security and privacy policies that constrain access to these data. In this work, we present a vision for privacy-preserving federated neural network architectures that permit data to remain at a custodian's institution while enabling the data to be discovered and used in neural network modeling. Using a diabetes dataset, we demonstrate that accuracy and processing efficiencies using federated deep learning architectures are equivalent to the models built on centralized datasets.

**Keywords.** federated learning, data decentralization, health data analytics, privacy preserving machine learning, neural network

## 1. Introduction

There is a treasure trove of knowledge locked inside the clinical data that pervades every modern health care organisation, ranging from electronic health records to bio-marker information found in genomic databases. The excitement around deep learning is real and finding ways to improve deep learning methods occupies the research agendas of scientists globally. However, the training of neural network models requires vast amounts of data, usually sourced by consolidating data from disparate systems to a single database.

Federated learning (FL) [1] can solve this problem by allowing a machine learning model to be trained without needing to consolidate the training data into a single location. FL has also been applied in the medical field for facilitating various distributed configurations on raw patient data using SplitNN [2] when the data is split in multiple hospitals. These studies shows that FL is a promising approach to obtain robust and safe models in terms of performance and data protection, although there are still open technical and health information sharing policies to be resolved before federated learning could be widely adopted in real world healthcare settings [3].

We assert a novel solution that lies in not attempting to consolidate data from these disparate systems but instead to employ a federated neural network approach. A federated approach eliminates the problems related to inadvertent data leakage,

---

<sup>1</sup> Corresponding Author, Helen Chen, School of Public Health and Health Systems, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada; E-mail: helen.chen@uwaterloo.ca

regulatory compliance, incompatible data schemas, etc. while garnering a large and rich dataset able to support deep learning. We demonstrate how this is possible even when needing to overcome: (i) real-world scenarios of how data might be split between host organisations, and (ii) performance issues. We validate our contribution by undertaking a real-world evaluation using the Pima Indian Diabetes dataset [4]. The Pima Indian diabetes dataset is publicly available and is extracted from the National Institute of Diabetes and Digestive and Kidney Diseases<sup>2</sup>. It has 768 rows and 8 columns. We acknowledge the very limited features and the small size of this dataset. The reason for using this dataset in this study is its simplicity, which allows us to perform quick comparison of performances of the different learning processes.

## 2. Federated Neural Network Architectures for Healthcare

In the continuum of care, a patient may receive care from multiple care providers, where patient data are collected and resides in siloed or centralized databases. There are three scenarios in which a patient's data may be stored. (i) **Horizontal partition** where the data matrix in the virtual central repository is horizontally partitioned, as shown in Figure 1(a). This partition corresponds to the scenario where different institutes use the same schema to collect and store patient data. (ii) **Vertical partition** where the data are vertically divided by features and the label may reside in one or multiple partitions as shown in Figure 1(b). Vertical partitioning of a dataset corresponds to the scenario where a patient's treatment occurs in multiple hospitals where data are collected in each care facility. (iii) **Mixed partition** which is a combination of vertical and horizontal splits, as shown in Figure 1(c). This pattern corresponds to the scenario where some patients' records are entirely stored in one institute, while other patients' records have been stored in different institutes.




Figure 1. Data Partitions

We explore two different federated neural network architectures as illustrated in Figure 2:

**Architecture 1:** This architecture applies to horizontally partitioned data where there are multiple data owners (e.g. hospitals) and a single server. The central server sends a model initialized with the same values to each data owner. Each data owner trains this model on its local data and sends the weights back to the server. The server receives weights from all participants and updates the central model with the averaged weights.

<sup>2</sup> <https://www.kaggle.com/uciml/pima-indians-diabetes-database>.

The central model then is sent back to each data owner for further training. The training process repeats until the model converges. This architecture requires that the data owned by each party has a full set of features.

**Architecture 2:** This architecture is inspired from the vertical split network from SplitNN [2]. A single neural network is partially trained separately at the data owners and then again at the server which keeps an updated schema with information about the owner of each feature. Each data owner calculates intermediary results of the network and communicates the results to the central server. The server trains the rest of the network based on the intermediary results and sends each owner their respective weights. The training process then repeats until convergence. This architecture allows the server to select data from wherever it resides and enables the server to train in both the vertical and mixed data split scenarios. In the mixed data scenario, data owners compute on partial data, treating missing information as empty values.

The neural network model used for the experiment is a linear network with one hidden layer. The activation function used for the layers is ReLU. The output of the model is passed through a softmax and loss is calculated by the negative log likelihood function. All models are trained 10 times and the results are averaged. All experiments are carried out on a computer with a core i5 3.4 GHz processor and 8GB RAM. The algorithm code is written in Python using PyTorch.




Figure 2. Proposed Federated Architectures

3. Experiments

We evaluated the accuracy of the models trained in the federated learning architectures and compared them with the centralized model with all data located at one place and the local model where each party trains on its own data.

**Experiment 1 (Accuracy equivalence):** We ran both architectures and the central model by initializing them firstly with separate weights and secondly with the same initial random weights. We empirically showed that our federated architectures are equivalent in Figure 3a. From Figure 3b, we can observe that the training accuracy of architecture 2 and the central model are almost equal and that Architecture 1 is approximating towards the same accuracy.

**Experiment 2 (Data equally split):** We distributed the whole diabetes dataset into 4 equal halves and use architecture 1 to study the horizontal splitting. We present the final testing accuracy of each hospital and the federated model in Table 1. Notably, the federated model achieves better final testing accuracy than the local models.

**Experiment 3 (Data unequally split):** In this experiment, the data was split with 30%, 30%, 35% and 5% data respectively saved in each hospital. The final testing accuracy of each hospital and the federated model is reported in Table 1. The testing accuracy of the hospitals with lesser amounts of data is at least 2% lower than the hospital which contains the largest amount of data. These hospitals can gain access to better models by participating with other hospitals that have access to more information.

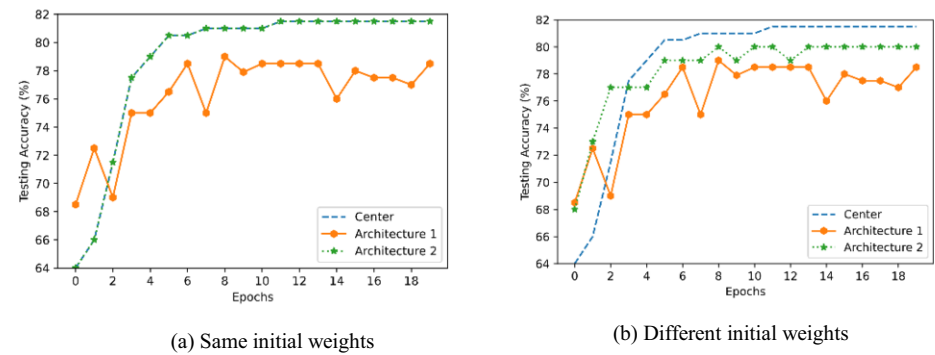


Figure 3. Experiment 1 testing accuracies

**Table 1.** Model accuracy comparison – FL model achieves best or equal accuracy compared to local models. In unequal split case, Hospital 4 has the least amount of training data and achieves worst accuracy

| Equal Split(%)         |             | Unequal Split(%)       |             |
|------------------------|-------------|------------------------|-------------|
| Hospital 1(25%)        | 74.8        | Hospital 1(30%)        | 76.1        |
| Hospital 2(25%)        | 73.1        | Hospital 2(30%)        | 73.5        |
| Hospital 3(25%)        | 74.0        | Hospital 3(35%)        | 77.9        |
| Hospital 4(25%)        | 74.0        | Hospital 4(5%)         | 70.9        |
| <b>Federated Model</b> | <b>75.3</b> | <b>Federated Model</b> | <b>77.9</b> |

4. Discussion and Future Work

FL architectures enable multiple parties to collaboratively construct a learning model. However, there are a number of issues that require special attention when developing a federated neural network in healthcare. (i) **Imbalanced class:** Imbalanced class is prevalent in healthcare datasets (e.g. the occurrence of a rare disease in a general population dataset is small). FL would fail in such a case as the local optimization would be highly skewed for a particular class, resulting in no convergence of the model [5]. In future work, we plan to modify the current federated architectures to enhance the learning for imbalanced distributed datasets. (ii) **Real time latency issues:** One limitation of the proposed FL architectures is its dependency on large volume of fast communication between the model host and the participating sites. Any communication failure may cause a slow or sub-optimal convergence of the model. Several works proposed enhancement of communication efficiency in FL by dropping a subset of data owners at each iteration [6]. However, this strategy needs to be evaluated in the vertical and mixed split settings. (iii) **Data privacy:** Privacy protection is of vital importance in health data sciences. FL models are prone to membership inference attacks [7]. To address these issues, a naïve solution is to deidentify personal data. However, instead of excluding

features from the dataset to protect privacy, privacy-preserving techniques, such as differential privacy [8] could be incorporated in the FL network. Finally, there are also other practical challenges for deploying a real-world federated learning architecture including but not limited to data provenance, geographic coverage or data quality.

## 5. Conclusion

In this paper, we presented FL architectures using horizontally, vertically and mixed-split data reflecting the real-world healthcare settings. We demonstrated the robustness of FL architecture using both equal and unequal data distribution scenarios. The models trained on the FL network achieved equivalent testing accuracy as the centralized models. This study highlights the feasibility of FL architectures in reaping the benefits of centralized data without the pain of creating centralized data for learning purposes.

**Acknowledgement:** This research was supported by Canada Mitacs Accelerate Program and F. Hoffmann-La Roche Ltd Canada (project #IT09798 “Clinical Analytics for Real-World Evidence Platform”).

## References

- [1] Yang Q, Liu Y, Chen T, and Tong Y. Federated machine learning: Concept and applications. In *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2019 10(2):1-19
- [2] Vepakomma P, Gupta O, Swedish T, and Raskar R. Split learning for health: Distributed deep learning without sharing raw patient data. *ICLR AI for Social Good Workshop* Available at: [https://aiforsocialgood.github.io/iclr2019/accepted/track1/pdfs/31\\_aisg\\_iclr2019.pdf](https://aiforsocialgood.github.io/iclr2019/accepted/track1/pdfs/31_aisg_iclr2019.pdf).
- [3] Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, Bonawitz K, Charles Z, Cormode G, Cummings R, et al. Advances and open problems in federated learning. 2019 arXiv preprint arXiv:1912.04977. Available at <https://arxiv.org/abs/1912.04977>.
- [4] Smith JW, Everhart JE, Dickson WC, Knowler WC, and Johannes RS. Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. In *Proc. Annu Symp Comput Appl Med Care, AMIA* 1988 Nov 9, p 261-65.
- [5] He H, Garcia EA. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 2009 21(9): p 1263-1284
- [6] Reisizadeh A, Mokhtari A, Hassani H, Jadbabaie A, and Pedarsani R. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *Proc. 23<sup>rd</sup> Int Conf Artificial Intelligence and Statistics*, 2020 PMLR 108: 2021-2031
- [7] Wang Z, Song M, Zhang Z, Song Y, Wang Q, and Qi H. Beyond inferring class representatives: User-level privacy leakage from federated learning. In *IEEE INFOCOM 2019 - IEEE Conf on Computer Communications*, Paris, France, 2019 p 2512-2520
- [8] Dwork C, McSherry F, Nissim K, and Smith A. Calibrating noise to sensitivity in private data analysis. Halevi S., Rabin T. (eds) *Theory of Cryptography. TCC 2006. Lecture Notes in Computer Science*, vol 3876. Springer, Berlin, Heidelberg, p. 265-84.