Differential Privacy-enabled Federated Learning for Sensitive Health Data

Olivia Choudhury[1, *], **Aris Gkoulalas-Divanis**[2], **Theodoros Salonidis**[3], **Issa Sylla**[1],
Yoonyoung Park[1], **Grace Hsu**[4, **], and **Amar Das**[1]

[1]IBM Research Cambridge, Cambridge, MA, USA
[2]IBM Watson, Cambridge, MA, USA
[3]IBM T.J. Watson Research Center, Yorktown Heights, NY, USA
[4]Massachusetts Institute of Technology, Cambridge, MA, USA
[*]Corresponding author: `olivia.choudhury1@ibm.com`
[**]Research done during internship at IBM Research Cambridge

Abstract

Leveraging real-world health data for machine learning tasks requires addressing many practical challenges, such as distributed data silos, privacy concerns with creating a centralized database from person-specific sensitive data, resource constraints for transferring and integrating data from multiple sites, and risk of a single point of failure. In this paper, we introduce a federated learning framework that can learn a global model from distributed health data held locally at different sites. The framework offers two levels of privacy protection. First, it does not move or share raw data across sites or with a centralized server during the model training process. Second, it uses a differential privacy mechanism to further protect the model from potential privacy attacks. We perform a comprehensive evaluation of our approach on two healthcare applications, using real-world electronic health data of 1 million patients. We demonstrate the feasibility and effectiveness of the federated learning framework in offering an elevated level of privacy and maintaining utility of the global model.

1 Introduction

Learning from real-world health data has proven effective in multiple healthcare applications, resulting in improved quality of care [1, 2, 3], generating medical image diagnostic tools [4, 5, 6], predicting disease risk factors [7, 8, 9], and analyzing genomic data for personalized medicine [10, 11, 12]. Healthcare data is often segregated across data silos, which limits its potential for insightful analytics. Data access is further restricted due to regulatory policies mandated by the US Health Insurance Portability and Accountability Act (HIPAA)[1] and EU General Data Protection Regulation (GDPR) [13]. Traditional or centralized machine learning algorithms require aggregating such distributed data into a central repository for the purpose of training a model. However, this incurs practical challenges, such as regulatory restrictions on sharing patient-level sensitive data, high resources required for transferring and aggregating the data, as well as a high risk associated with introducing a single point of failure. Leveraging such data while complying with data protection policies requires re-thinking data analytics methods for healthcare applications.

Federated learning (FL) offers a new paradigm for training a global machine learning model from data distributed across multiple data silos, eliminating the need for raw data sharing [14]. Once a global model is shared across sites, each site trains the model based on its local data. The parameter updates of the local models are subsequently sent to an aggregation server and incorporated into the global

---

[1]https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html

model. This is repeated until a convergence criterion of the global model is satisfied. The merit of FL has been recently demonstrated in several real-world applications, including image classification [15] and language modeling [14]. It is especially relevant in healthcare applications, where data is rife with personal, highly-sensitive information, and data analysis methods must comply with regulatory requirements [16].

In this paper, we propose, implement, and evaluate a FL framework for analyzing distributed healthcare data. FL provides a first level of privacy protection by training a global model without sharing raw data among sites. However, in certain cases, FL may be vulnerable to inference attacks [17, 18, 19]. Although state-of-the-art approaches to address such attacks are based on differential privacy [20], their performance has not been investigated for healthcare applications. To this end, we extend the FL framework with a distributed differential privacy mechanism and investigate its performance.

We perform a comprehensive empirical evaluation using two real-world health datasets comprising electronic health records and administrative claims data of 1 million patients. We analyze the effect of different levels of privacy on the performance of the FL model and derive insights about the effectiveness of FL with and without differential privacy in healthcare applications. More specifically, we show that FL without differential privacy can provide model performance close to the hypothetical case where all data is centralized. We then study how ϵ-differential privacy impacts the performance of the produced global FL model for a given level of privacy. The results show that although differential privacy offers a strong level of privacy, it deteriorates the predictive capability of the produced global models due to the excessive amount of noise added during the distributed FL training process.

2   Methods

2.1   Federated learning

For the purpose of illustration, we consider classification algorithms that are amenable to gradient descent optimization. For a general binary classification problem, let the features, denoted by x_k (for the k^{th} feature), be drawn from a feature space $\mathbb{X}$. The corresponding labels y_k are drawn from the label space $\mathbb{Y} := \{-1, 1\}$. Let the features corresponding to positive labels be denoted by $\mathbb{X}_+$ and those corresponding to negative labels by $\mathbb{X}_-$, that is

$$\mathbb{X}_+ = \{x_k \in \mathbb{X} : y_k = +1\} \quad \text{and} \quad \mathbb{X}_- = \{x_k \in \mathbb{X} : y_k = -1\}.$$

For any $x_k^+ \in \mathbb{X}_+$ and $x_k^- \in \mathbb{X}_-$, the objective of this classification is to construct a function $f : \mathbb{X} \to \mathbb{Y}$ such that

$$f(x_k^+) = +1 \quad \text{and} \quad f(x_k^-) = -1.$$

As discussed later, for our use cases of predicting adverse drug reaction and in-hospital mortality, we denote cases of ADR and in-hospital mortality as labels $y_k = +1$, and cases of non-ADR and non-mortality as $y_k = -1$. For a FL setup with N sites, let $\mathcal{D}_i$, with feature set $\{\mathbb{X}_{\text{train}}^i\}_{i=1}^N$ and corresponding label set $\{\mathbb{Y}_{\text{train}}^i\}_{i=1}^N$ be the local training data at site i, where $i \in N$. Based on the use case, a global model is shared with each site, which trains the model on its local data $\mathcal{D}_i$, where $i \in N$. During local model training, based on the given learning rate, number of epochs, and batch, we compute average gradient $(\triangledown F_i(w))$ with respect to its current model parameter w. We then compute weighted average to aggregate the parameter updates from the local models. The process is repeated until a convergence criterion, such as minimization of loss function, is satisfied. For further details on implementing a FL model for healthcare applications, we refer the readers to [21].

2.2   Differential Privacy

Differential privacy [22, 23, 24] is a widely-used standard for privacy guarantee of algorithms operating on aggregated data. A randomized algorithm $\mathcal{A}(\mathcal{D})$ satisfies ϵ-differential privacy if for all datasets $\mathcal{D}$ and $\mathcal{D}'$, that differ by a single record, and for all sets $\mathcal{S} \in \mathcal{R}$, where $\mathcal{R}$ is the range of $\mathcal{A}$,

$$Pr[\mathcal{A}(\mathcal{D}) \in \mathcal{S}] \le e^\epsilon Pr[\mathcal{A}(\mathcal{D}') \in \mathcal{S}].$$

where ϵ, a privacy parameter, is a non-negative numbers. This implies that any single record in the dataset does not have a significant impact on the output distribution of the algorithm.

There are several methods for generating an approximation of $\mathcal{A}$ that satisfies differential privacy. Based on the different approaches of adding noise, they can be categorized into input perturbation, output perturbation, exponential mechanism, and objective perturbation [25]. Existing work on the application of differential privacy in machine learning have primarily focused on models that are trained on a centralized dataset [26, 27, 25]. As a recent advancement in privacy-preserving FL, the authors in [19] adopted output perturbation or *sensitivity method*. However, prior research has shown the effectiveness of objective perturbation, with theoretical guarantee, in outperforming the output perturbation approach [28]. Hence, in this work, we explore the potential of differential privacy, based on objective perturbation, in the context of FL. For the task of classification, we add noise to the objective function of the optimization to obtain a differentially private approximation. At each site, the noise is added to the objective function of the model to produce a minimizer of the perturbed objective. For a comprehensive background on differential privacy with objective perturbation, we refer the readers to [29, 28].

3 Evaluation

3.1 Use cases and data preparation

Developing FL models in a privacy-preserving manner is very important, especially in the context of healthcare, where patient data are extremely sensitive. To evaluate our proposed approach, we consider two major tasks for improving the health outcome of patients: (a) prediction of adverse drug reaction (ADR), and (b) prediction of mortality rate. ADR is a major cause of concern amongst medical practitioners, pharmaceutical industry, and the healthcare system[3]. As healthcare data is distributed across data silos, obtaining a sufficiently large dataset to detect such rare events poses a challenge for centralized learning models. For the purpose of ADR prediction, we used Limited MarketScan Explorys Claims-EMR Data (LCED), which comprises administrative claims and electronic health records (EHRs) of over 1 million commercially insured patients. It consists of patient-level sensitive features, such as demographics, habits, diagnosis codes, outpatient prescription fills, laboratory results, and inpatient admission records. We selected patients who received a nonsteroidal anti-inflammatory drug (NSAID) to predict the development of peptic ulcer disease following the initiation of the drug. The selected cohort comprised $921,167$ samples.

For the second use case, we considered the task of modeling in-hospital patient mortality. An accurate and timely prediction of this outcome, particularly for patients admitted to an intensive care unit (ICU), can significantly improve quality of care. For this task, we used the Medical Information Mart for Intensive Care (MIMIC III) data [30]. MIMIC III is a benchmark data set, from where we derived multivariate time series from over $40,000$ ICU stays and labels to model mortality rate. As discussed in [31], we selected 17 physiological variables, including demographic details, each comprising 6 different sample statistic features on 7 different subsequences of a given time series, resulting in 714 features per times series. The cohort consisted of $21,139$ ICU stays.

3.2 Experimental setup

For the tasks of ADR and mortality prediction, we used three classification algorithms, amenable to distributed solution using gradient descent, namely perceptron, support vector machine (SVM), and logistic regression. To evaluate the models, prior to and after employing privacy-preserving mechanism, we measure their utility in terms of F1 score. The models were trained on 70% of the data with 5-fold cross-validation. We considered 10 sites for the federated setup. All experiments were executed on an Intel(R) Xeon(R) E5-2683 v4 2.10 GHz CPU equipped with 16 cores and 64 GB of RAM. The results were reported following 10 rounds of iteration.

3.3 Comparative analysis

To establish benchmark results, we first evaluated the performance of the centralized learning models for the tasks of predicting ADR and in-hospital mortality. We then analyzed the performance of FL models trained on distributed data. For ϵ-differential privacy, we measured the privacy-utility trade-off for a given range of the privacy parameter ϵ. Figure 1 (a) and (b) present the utility, measured

---

[3]https://www.fda.gov/drugs/informationondrugs/ucm135151.htm

by F1 score, when differential privacy is applied using typical values of parameter $\epsilon \in [0.01, 0.5]$, for LCED and MIMIC data, respectively. As ϵ increases, the level of privacy degrades, thereby improving the utility or predictive capability of the models. This is consistent across all three classification algorithms and for both datasets.
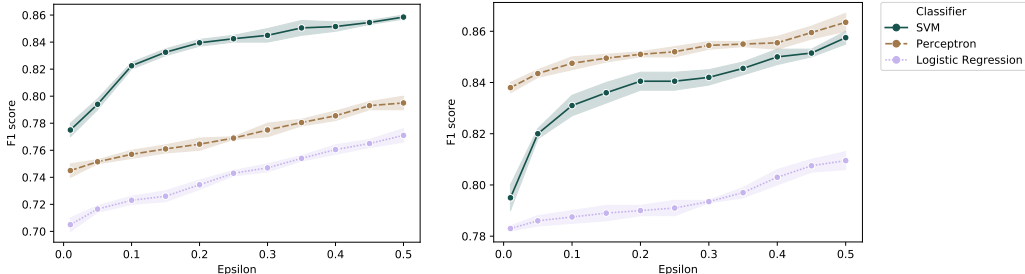

Figure 1: Effect of varying ϵ in ϵ-differential privacy for (a) ADR prediction using LCED, and (b) in-hospital mortality prediction using MIMIC data.
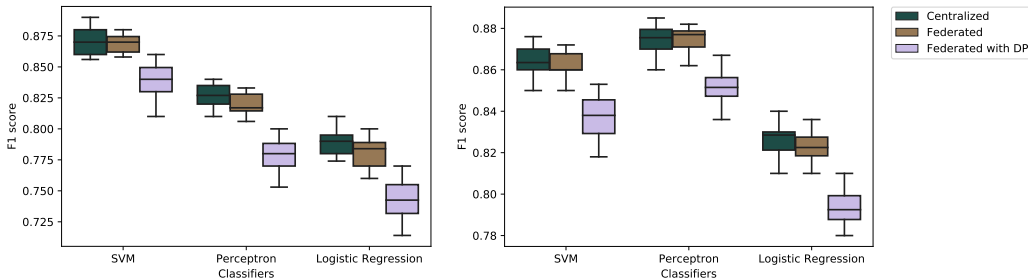

Figure 2: Comparison of F1 score with (a) LCED, and (b) MIMIC data, between centralized learning, FL, and FL with ϵ-differential privacy (Federated with DP), where $\epsilon = [0.01, 0.5]$.

We also compare and contrast the performance of centralized learning, FL, and FL with ϵ-differential privacy in terms of utility, for $\epsilon \in [0.01, 0.5]$. This indicates the level of utility that can be attained for an acceptable range of the privacy parameter ϵ. As shown in Figure 2 (a) and (b), FL achieves comparable performance to centralized learning for both tasks of ADR and mortality prediction. Although differential privacy guarantees a given level of privacy, as set by parameter ϵ, it leads to a significant deterioration of the utility of the federated model.

It must be noted that existing studies advocating the use of differential privacy in a FL setup have not been performed on real data. Moreover, as discussed in [19], the model performance can only be preserved for a very large number of sites, in the order of 1000, but takes a severe hit in the case of fewer sites. Such an assumption of large-scale setup is not realistic for healthcare applications, where sites are typically hospitals or providers, and each site may not have sufficient data for deep learning models to be applicable. This drives the need to leverage data from other sites for constructing more accurate models. Hence, although differential privacy is widely adopted for preserving the privacy of machine learning models, it can yield lower utility in a FL setup, particularly in real-world healthcare applications. This necessitates exploring alternative privacy-preserving approaches that can achieve high model performance while offering sufficient privacy [32].

4 Conclusion

The availability of electronic health data poses several opportunities to leverage machine learning for deriving insightful data analytics. In this paper, we implemented a federated learning approach to mitigate the challenges associated with centralized learning. We further explored the potential of differential privacy in preserving the privacy of federated learning models, while maintaining high model performance in the context of real-world health applications. Through experimental evaluation, we show that although differential privacy is being readily adopted in a federated setup, it can lead to a significant loss in model performance for healthcare applications. This necessitates the research and proposal of alternative approaches towards offering privacy in federated learning for healthcare applications.

References

[1] E. Gultepe, J. P. Green, H. Nguyen, J. Adams, T. Albertson, and I. Tagkopoulos, "From vital signs to clinical outcomes for patients with sepsis: a machine learning basis for a clinical decision support system," *Journal of the American Medical Informatics Association*, vol. 21, no. 2, pp. 315–325, 2013.

[2] A. E. Johnson, M. M. Ghassemi, S. Nemati, K. E. Niehaus, D. A. Clifton, and G. D. Clifford, "Machine learning and decision support in critical care," *Proceedings of the IEEE. Institute of Electrical and Electronics Engineers*, vol. 104, no. 2, p. 444, 2016.

[3] A. Chakrabarty, S. Zavitsanou, T. Sowrirajan, F. J. Doyle III, and E. Dassau, "Getting IoT-ready: The face of next generation artificial pancreas systems," in *The Artificial Pancreas*. Elsevier, 2019, pp. 29–57.

[4] B. J. Erickson, P. Korfiatis, Z. Akkus, and T. L. Kline, "Machine learning for medical imaging," *Radiographics*, vol. 37, no. 2, pp. 505–515, 2017.

[5] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest *et al.*, "The multimodal brain tumor image segmentation benchmark (BRATS)," *IEEE transactions on medical imaging*, vol. 34, no. 10, pp. 1993–2024, 2014.

[6] E. Gibson, W. Li, C. Sudre, L. Fidon, D. I. Shakir, G. Wang, Z. Eaton-Rosen, R. Gray, T. Doel, Y. Hu *et al.*, "NiftyNet: a deep-learning platform for medical imaging," *Computer methods and programs in biomedicine*, vol. 158, pp. 113–122, 2018.

[7] R. Poplin, A. V. Varadarajan, K. Blumer, Y. Liu, M. V. McConnell, G. S. Corrado, L. Peng, and D. R. Webster, "Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning," *Nature Biomedical Engineering*, vol. 2, no. 3, p. 158, 2018.

[8] T. Zheng, W. Xie, L. Xu, X. He, Y. Zhang, M. You, G. Yang, and Y. Chen, "A machine learning-based framework to identify type 2 diabetes through electronic health records," *International journal of medical informatics*, vol. 97, pp. 120–127, 2017.

[9] K. Kourou, T. P. Exarchos, K. P. Exarchos, M. V. Karamouzis, and D. I. Fotiadis, "Machine learning applications in cancer prognosis and prediction," *Computational and structural biotechnology journal*, vol. 13, pp. 8–17, 2015.

[10] M. W. Libbrecht and W. S. Noble, "Machine learning applications in genetics and genomics," *Nature Reviews Genetics*, vol. 16, no. 6, p. 321, 2015.

[11] O. Choudhury, A. Chakrabarty, and S. J. Emrich, "Hecil: A hybrid error correction algorithm for long reads with iterative learning," *Scientific reports*, vol. 8, no. 1, p. 9936, 2018.

[12] A. L. Beam and I. S. Kohane, "Big data and machine learning in health care," *Jama*, vol. 319, no. 13, pp. 1317–1318, 2018.

[13] Agencia Espanola Proteccion Datos, "K-Anonymity as a Privacy Measure," 2019, https://www.aepd.es/media/notas-tecnicas/nota-tecnica-kanonimidad-en.pdf.

[14] H. B. McMahan, E. Moore, D. Ramage, S. Hampson *et al.*, "Communication-efficient learning of deep networks from decentralized data," *arXiv preprint arXiv:1602.05629*, 2016.

[15] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1205–1221, 2019.

[16] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, "Federated learning of predictive models from federated Electronic Health Records," *International journal of medical informatics*, vol. 112, pp. 59–67, 2018.

[17] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," *arXiv preprint arXiv:1807.00459*, 2018.

[18] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2017, pp. 1175–1191.

[19] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.

[20] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*.   Springer, 2006, pp. 486–503.

[21] O. Choudhury, Y. Park, T. Salonidis, A. Gkoulalas-Divanis, I. Sylla, and A. Das, "Predicting Adverse Drug Reactions on Distributed Health Data using Federated Learning," *American Medical Informatics Association (AMIA)*, 2019.

[22] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*.   Springer, 2006, pp. 265–284.

[23] C. Dwork, "A firm foundation for private data analysis," *Communications of the ACM*, vol. 54, no. 1, pp. 86–95, 2011.

[24] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[25] A. D. Sarwate and K. Chaudhuri, "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data," *IEEE signal processing magazine*, vol. 30, no. 5, pp. 86–94, 2013.

[26] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*.   ACM, 2016, pp. 308–318.

[27] Z. Ji, Z. C. Lipton, and C. Elkan, "Differential privacy and machine learning: a survey and review," *arXiv preprint arXiv:1412.7584*, 2014.

[28] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, "Differentially private empirical risk minimization," *Journal of Machine Learning Research*, vol. 12, no. Mar, pp. 1069–1109, 2011.

[29] K. Chaudhuri and C. Monteleoni, "Privacy-preserving logistic regression," in *Advances in neural information processing systems*, 2009, pp. 289–296.

[30] A. E. Johnson, T. J. Pollard, L. Shen, H. L. Li-wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "MIMIC-III, a freely accessible critical care database," *Scientific data*, vol. 3, p. 160035, 2016.

[31] H. Harutyunyan, H. Khachatrian, D. C. Kale, G. Ver Steeg, and A. Galstyan, "Multitask learning and benchmarking with clinical time series data," *Scientific data*, vol. 6, no. 1, p. 96, 2019.

[32] O. Choudhury, A. Gkoulalas-Divanis, T. Salonidis, I. Sylla, Y. Park, G. Hsu, and A. Das, "Anonymizing data for privacy-preserving federated learning," *arXiv preprint arXiv:2002.09096*, 2020.