Sandia National Laboratories

Federated Learning and Differential Privacy: What might AI-Enhanced co-design of microelectronics learn?

Evercita C. Eugenio

May 2022

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

ABSTRACT

Data is a valuable commodity, and it is often dispersed over multiple entities. Sharing data or models created from the data is not simple due to concerns regarding security, privacy, ownership, and model inversion. This limitation in sharing can hinder model training and development. Federated learning can enable data or model sharing across multiple entities that control local data without having to share or exchange the data themselves. Differential privacy is a conceptual framework that brings strong mathematical guarantee for privacy protection and helps provide a quantifiable privacy guarantee to any data or models shared. The concepts of federated learning and differential privacy are introduced along with possible connections. Lastly, some open discussion topics on how federated learning and differential privacy can tied to AI-Enhanced co-design of microelectronics are highlighted.

KEY-TAKEAWAYS

- Federated learning can enable model or data sharing across multiple entities that hold or control local data without having to actually share or exchange the data themselves.
- Differential privacy can provide a quantifiable provable guarantee to amount of privacy leaked by models or datasets being shared.
- Together federated learning and differential privacy could address some of the open discussion topics in AI-Enhanced co-design of microelectronics.

INTRODUCTION

Data privacy in today's world is a hot-button topic. Companies such as Netflix, Facebook, Microsoft, Google, among others have caused an uproar in how they collect data and how they share data. Headlines in newspapers, magazines, and journals around the world, highlight the growing concerns with data privacy, data collection, data protection and data sharing. These concerns are ever growing and serve as motivation for work in the privacy field.

While privacy breaches continue to be a huge concern, it has not stopped the requests for data sharing or data release. Government agencies, business, survey and research organizations, and medical institutions are constantly being asked to release and share more and more of their data for transparency and accountability. Thus, handling all this data in a way that protects the confidentiality of the data subjects' identities and sensitive attributes while maintaining the statistical usability and accuracy of the data set has developed into a critical area of study.

One of the most common ways to "protect" data is to simply anonymize the data (i.e. remove identifying information or sensitive characteristics). However, simply anonymizing data does not simply solve the problems with record linkage. For example, if there is a data set that has name, address, phone number and education. And another data set with medical diagnoses, procedures, and payment information. All you would need is gender, birth date and zip code to link the two data sets. Sweeny 2013 was able to identify 130 of the people in the personal genome project simply by using voter data.

FEDERATED LEARNING

Federated learning is a technique in machine learning that enables model training over multiple entities (these entities could be servers, devices, databases, etc.) that hold local data without actually exchanging the data themselves (Kone"c" ˇ"n" "y" ´ 2015). Thus, even though one may not have access to all the data one can still do model training and get machine learning models based on all the data. This is a powerful tool since a machine learning model is only as good as the data it is built on, therefore having a larger training dataset is beneficial. Further, data collection is often difficult, and organizations are limited based on the cost, time, and energy required to collect data. This often leads to multiple agencies to collecting data of similar nature, leaving data disjoint with many owners. Federated learning could help facilitate sharing in these instances.

Consider the example where there are 7 datasets and each one is owned by a different entity. Ideally what all 7 datasets would be combined together to one massive dataset to run analysis on. However, that is not very practical. In fact, most people are not willing to simply share their data so that it can be combined with other data. Complications and concerns regarding privacy, security, ownership and model inversion further constrain the sharing, and often limit model training and development.

There are three major types of federated learning: 1) centralized federated learning; 2) decentralized federated learning; and 3) heterogenous federated learning. Figure 1 provides an example of a centralized federated learning system. In this system there are multiple entities that all own some dataset. The central server, a trusted entity, sends some an initial model to all the entities, and requests them to get a local model based on the data they own. These local models are then shared back to the central server which can then combine the multiple local models into a global model. This global model is now based on all the data owned by multiple entities, but the central server never had to access the data points themselves. Multiple exchanges between the central server and multiple entities may be required to get a more accurate global model.

For decentralized federated learning, there is no central server, and the nodes are able to coordinate themselves to obtain the global model. This setup prevents single point failures as the model updates are exchanged only between interconnected nodes without requiring a central server to orchestrate the entire process.

Heterogenous federated learning has become a more commonly used type of federated learning system. This is because unlike many of the existing federated learning methods, this does not assume that the local models all share the same global architecture. HetroFL from Diao et al. 2020 proposed a new framework that allowed training of heterogenous local models with varying computation complexities to still get a single global model.

Lastly in federated learning it is important to consider whether the data hosted by each of the entities or nodes is balance or not. That must be considered to create the global models. Thus overall, federated learning is often employed, as it enables training an algorithm over multiple data sources without exchanging the data samples themselves. Unfortunately, this current system for sharing models is marred by privacy concerns regarding reverse engineering the original data based on the released model parameters.

DIFFERENTIAL PRIVACY

Differential privacy is a conceptual framework that brings strong mathematical guarantee for privacy protection and helps provide a quantifiable privacy guarantee to any data or models shared. Differential privacy ensures that the addition or removal of a single database item does not substantially affect the outcome of any analysis. This means that just because a single person is or isn't in a database, you should not be able to deduce any information about them. More formally, the definition of differential privacy is as follows. Let $\mathcal{K}$, a mechanism to be defined later, and let D_1 and D_2 be two databases that differ in at most one element. D_1 and D_2 differing by one individual and can be interpreted in two ways: 1) D_1 is one individual more or less than D_2, or D_1 and D_2 are of the same size but have difference in attributes values in exactly one individual.

> Definition (Dwork 2006): A randomized function $\mathcal{K}$ gives ϵ-differential privacy if for all datasets D_1 and D_2 be two databases that differ on at most one element, and all $S \subseteq Range(\mathcal{K}), Pr(\mathcal{K}(D_1) \in S) \leq e^{\epsilon} \times Pr(\mathcal{K}(D_2) \in S)$.

The formulation of privacy via differential privacy is robust and guards against the worst-case scenario as it does not impose any assumptions about the behavior or the background knowledge of data intruders. Thus, ϵ which is often referred to as the privacy budget and is pre-specified, is the amount of privacy used or information leaked. The larger ϵ is then the more information is leaked or less privacy, and the smaller ϵ is then the less information is leaked and the more privacy.

Differential privacy requires some noise to be added to queries of interest in order to protect privacy. Examples of noise adding mechanisms include the Laplace mechanism (Dwork 2006), Exponential Mechanism (McSherry and Talwar 2007), Gaussian mechanism (Dwork and Roth 2014), and the median mechanism (Roth and Roughgarden 2010). Additionally, noise adding mechanisms have been developed for specific statistical analyses such as contingency tables (Barak et al. 2007), principal component analysis (Chaudhuri et al. 2012), location privacy (Xiao and Xiong 2015) and graphs and social networks (Yan et al. 2016).

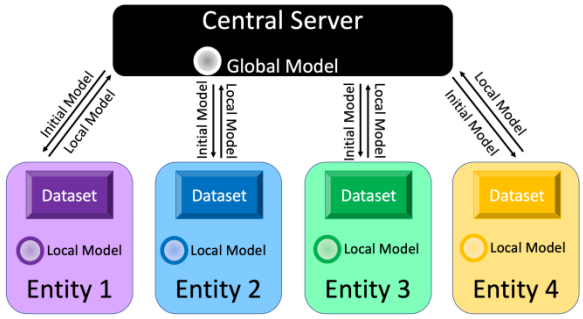

Figure 1: Centralized Federated Learning

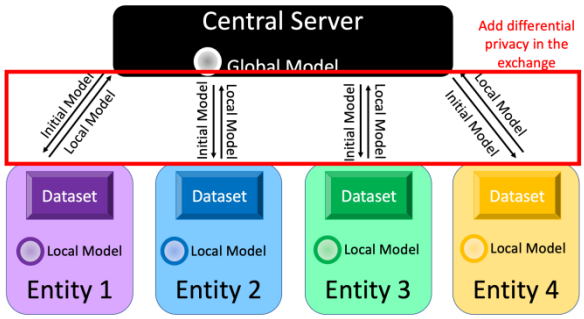

Figure 2: Federated Learning and Differential Privacy Together

FEDERATED LEARNING AND DIFFERENTIAL PRIVACY TOGETHER

Federated learning and differential privacy could be used together to help solve the problems surrounding data/model sharing and privacy. This is a growing area and there are many methods being developed to utilize both to facilitate sharing while still protecting privacy. Figure 2 provides a simple example of how these two concepts can used together. Specifically, prior to the local models being shared with the central server to create a global model, differential privacy could be used to add noise. This would allow for sharing of data without concerns about reverse engineering or leaking sensitive information.

OPEN DISCUSSION TOPICS

Federated learning and differential privacy could be utilized can tied to AI-Enhanced co-design of microelectronics. Some examples of open discussion topics include the following:

- **Data is valuable for training**: In the microelectronics application area, data may be limited or restricted, so having access to additional data would be extremely valuable.
- **Challenges to accessing data**: Accessing dataset related to microelectronics may be limited. If this is the case, then federated learning may prove to be a useful method to facilitate sharing while still limiting access to private data.
- **Importance of Privacy**: In the application area of microelectronics, there are instances where data cannot be shared due to the sensitive nature of the information in the datasets. Therefore, incorporating something like differential privacy could help ease these concerns and provide a quantifiable privacy value.
- **Possible use cases**: There are many possible use cases within microelectronics that could use either federated learning or differential privacy or both together.

REFERENCES

[1] Barak, B., Chaudhuri, K., Dwork, C., Kale, S., McSherry, F. and Talwar, K. (2007). Privacy, accuracy, and consistency too: a holistic solution to contingency table release. Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, pages 273-282.

[2] Chaudhuri, K., Sarwate, A., and Sinha, K. Near-optimal differentially private principal components. Advances in Neural Information Processing Systems, pages 989-997.

[3] Diao, E., Ding, J., and Tarokh, V. (2020). Heterofl: Computation and communication efficient federated learning for heterogeneous clients. arXiv preprint arXiv:2010.01264.

[4] Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In \textit{Proceedings of the Third Conference on Theory of Cryptography}, TCC'06, pages 265-284, Berlin, Heidelberg. Springer-Verlag.

[5] Dwork, C., and Roth, A. (2014). The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci., 9(3-4), 211-407.

[6] Konečný, J., McMahan, B., and Ramage, D. (2015). Federated optimization: Distributed optimization beyond the datacenter. arXiv preprint arXiv:1511.03575.

[7] McSherry, F. and Talwar, K. (2007). Mechanism design via differential privacy. In Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS '07, pages 94-103, Washington, DC, USA. IEEE Computer Society.

[8] Roth, A.and Roughgarden, T. (2010). Interactive privacy via the median mechanism. In Proceedings of the forty-second ACM symposium on Theory of computing, pages 765-774.

[9] Sweeney, L., Abu, A., and Winn, J. (2013). Identifying participants in the personal genome project by name (a re-identification experiment). arXiv preprint arXiv:1304.7605.

[10] Xiao, Y. and Xiong, L. (2015) Protecting locations with differential privacy under temporal correlations. Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pages 1298-1309.

[11] Yan, S., Pan, S., Zhao, Y., and Zhu, W.-T. (2016) Towards Privacy-Preserving Data Mining in Online Social Networks: Distance-Grained and Item-Grained Differential Privacy. Australasian Conference on Information Security and Privacy, pages 141-157.