

Federated Learning for Privacy-Preserving Medical Data Sharing in Drug Development

Mingxuan Yang^{1,a,*†}, Decheng Huang^{2,†}, Weixiang Wan³, Meizhizi Jin⁴

¹*Innovation Management and Entrepreneurship, Brown University, RI, USA*

²*Chemical and Biomolecular Engineering, University of Pennsylvania, Philadelphia, PA, USA*

³*Electronics & Communication Engineering, University of Electronic Science and Technology of China, Chengdu, China*

⁴*Management Information Systems, New York University, NY, USA*

a. rexcarry036@gmail.com

**corresponding author*

†These authors contributed equally to this work and should be considered co-first authors.

Abstract: This study explores the potential of Federated Learning (FL) to facilitate the sharing and collaboration of medical data in drug development under the premise of privacy protection. This paper systematically describes the core mechanism of federated learning, including the key technologies such as model parameter updating, differential privacy and homomorphic encryption, and their applications in drug development and medical data processing. Examples, such as NVIDIA Clara's Federated learning application and COVID-19 resource prediction, show that federated learning improves the efficiency of multi-party collaboration and model performance while ensuring data privacy, especially in areas such as finance and insurance, where data privacy is critical.

Keywords: Federated Learning, Data Privacy, Drug Development, Distributed AI.

1. Introduction

For a long time in the past, AI has been hailed as an important part of the industrial revolution, and continues to penetrate other industries, such as education, business, finance, manufacturing, as well as social media platforms and healthcare. With the continuous improvement of the data age and the emergence of advanced computer algorithms, people have a better opportunity to build new artificial intelligence models, and use it to achieve faster computing methods, so as to get more convenience. However, especially in healthcare, the centralization of a lot of data and AI faces multiple potential challenges in terms of privacy and regulation.

Hypothetically, if we can find a more efficient way to integrate AI of data into one and be able to break through the existing challenges while optimizing these risks, this will be a whole new area of research. [1] Federal Learning (FL) is the solution. Medical data is often scattered across different systems, and security and privacy concerns complicate its effective use. However, advances in AI have brought opportunities for integration and collaboration to this fragmented data. However, data is often scattered across siloed systems, and security and privacy concerns complicate its effective use.

Federated Learning (FL) is therefore the ideal solution to this problem. It allows data to remain local while fostering collaboration between agencies to build more robust AI models together without sacrificing data privacy and security. Through this approach, organizations can share information while protecting sensitive medical data, creating more possibilities for data utilization in drug development and driving innovation in medicine with privacy protection.

2. Related Work

To truly address these issues, we need to innovate traditional approaches and find more effective ways to decentralize data collection through data management tools and policies, while maintaining data security and privacy. In this process, it is not only necessary to continuously enrich AI technology[2], but also to continuously design how to translate data into more meaningful decision-making processes that drive progress across the industry. This concept is particularly relevant in drug development.

2.1. Drug development and federal learning

Drug development often relies on large amounts of patient data for research, with large amounts of medical data analyzed and validated at every step, from early drug discovery and preclinical trials to eventual clinical trials. However, due to the highly sensitive nature of patient data, how to share and utilize this data while ensuring privacy protection has been a challenge in the industry [3]. Traditional centralized data processing methods may not only lead to privacy leakage, but also face the problem of data silos, which hinders the effectiveness of cross-institutional cooperation and data sharing.

Federated Learning offers a possible solution to this problem. With federated learning, healthcare organizations can collaborate on training AI models based on a decentralized architecture while keeping data local. This means that organizations do not have to share raw patient data directly but can collaborate on AI model learning to improve the overall performance of the model. This approach can not only accelerate the research process of drug development, but also ensure the privacy and security of data and avoid violating relevant laws and regulations on patient privacy.

For example, at the clinical trial stage, federal learning could allow multiple healthcare institutions to jointly develop drug response prediction models without sharing raw data, screen potential drug candidates in a more efficient and precise way, and test drug efficacy and safety in different patient populations[4]. This will not only improve the efficiency of the trial, but also allow the patient data participating in the trial to maintain a high degree of privacy, laying a solid foundation for promoting personalized medicine and precision medicine.

To sum up, federal learning effectively promotes the sharing and utilization of medical data on the premise of protecting patient privacy, especially in drug development. [5]This innovative approach not only breaks down data silos and optimizes the research process, but also provides more reliable and efficient data support for decision making in drug development, thereby driving innovation and progress across the healthcare industry.

2.2. Key features of federated learning include

Data privacy protection: All sensitive data remains local to the data owner and does not leave its original location.

Distributed computing: Each participant independently trains the model using their local data and shares only the parameter updates or gradients of the model, not the original data.

Model collaborative optimization: Build a globally optimized model by aggregating model updates from different participants

This offers good prospects for the development of various AI models, especially where data can be easily accessed centrally to train these models[6]. However, in some fields, notably medicine, the centralization of data is often made difficult or even impossible by privacy, regulatory, competitive, and budgetary constraints. Figure 1 shows the architecture of a federated learning system as a means of distributed learning that can effectively mitigate the data deficiency challenge required to train AI models. By keeping data local and enabling collaboration across institutions, patient types, and countries, Federated Learning (FL) provides a solution for building robust AI models. At the same time, information management - getting the right information to the right people in a timely manner to support effective decision-making - is a pervasive problem.

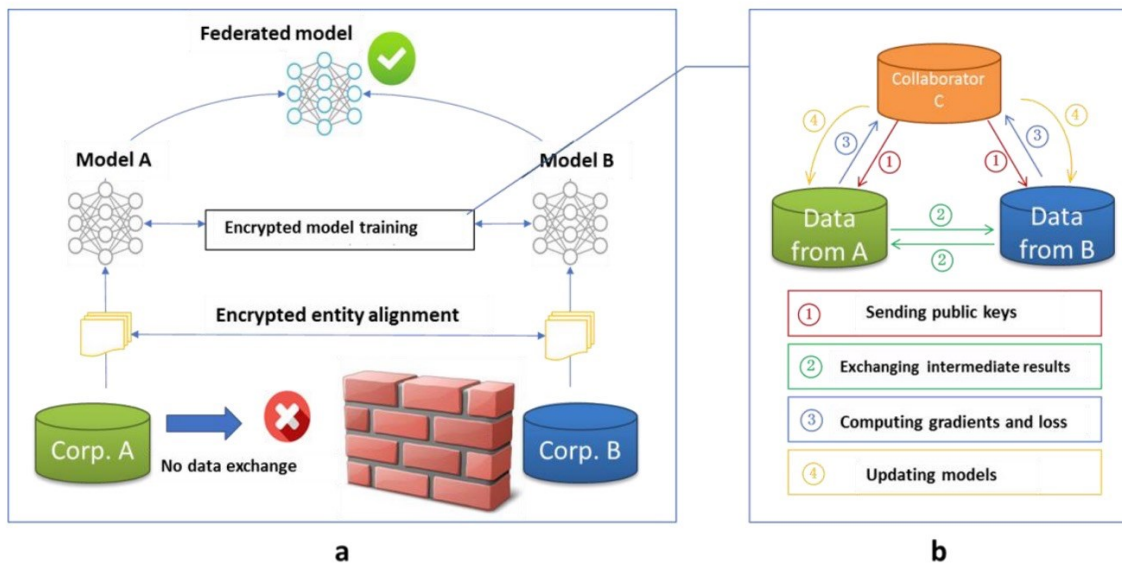


Figure 1: Architecture for a federated learning system

Data is often scattered in isolated, noninteroperable silos making it difficult to find and access. A mixture of structured and unstructured data, including text, audio, video, and imagery, with disparate formats and standards, is common. Meaningful information is embedded in massive amounts of irrelevant noise. Complicating matters are security and privacy concerns, cybersecurity threats, bandwidth limitations, platform challenges, and constrained budgets[7]. Once accessed, there must be a means to harness data; analyze, transform, and visualize it; and leverage it to execute automated processes and facilitate decision-making.

Adding to the challenge, there is no universal solution for every need and application. Even if there were, not every platform could simultaneously implement it due to cost and integration complexity. A steady state would never be reached due to ongoing modernization cycles that stretch over decades.

2.3. Federal learning integration for patient data privacy

One solution to this problem set is to integrate federated learning in the development of advanced models such as Large Language Model (LLM) [8] or other GenAI capability, infused with data analysis, automation, and decision support tools. Being federated, it accesses distributed non-standard data sources and integrates their data into a whole for processing and analysis. This is done without consolidation that stresses bandwidth and could compromise security and privacy. The federated learning system has been proven in finance for evaluating risk and detecting fraud, in transportation for route optimization and smart city design, and in healthcare to process patients' data and provide diagnosis analysis.

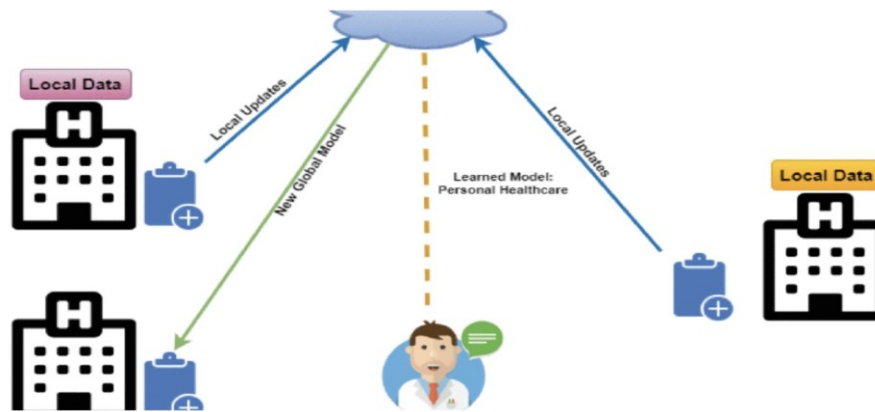


Figure 2: Federated Learning Framework for Healthcare

In healthcare, federated learning combined with the application of generative AI models (GenAI) is changing the way data is used and analyzed. Federated learning allows access to large-scale structured and unstructured data, such as medical textbooks, journal articles, research reports, electronic health records (EHRs), medical specialty exam questions and answers, audio, video, genomic data, and medical images, which are often distributed across multiple different data sources. Through the power of the GenAI[9] model, this fragmented data can be integrated and turned into actionable insights to help improve diagnostic accuracy, treatment outcomes, and patient health outcomes. This is all based on strict information security compliance requirements, ensuring that the data remains highly private and secure during use.

In drug development, the introduction of federal learning models is particularly important, especially for sharing medical data while protecting patient privacy. Drug development often requires multi-agency collaboration and large-scale data analysis, but traditional centralized data processing faces privacy breaches and legal and regulatory barriers. With federated learning, drug discovery teams can collaborate to train models on a global scale without transferring raw data, leveraging distributed data from different hospitals and research institutions.

Generative AI models can extract effective information from these diverse medical data and provide accurate drug response prediction and personalized treatment plans. This approach will not only significantly accelerate the development of new drugs, but also ensure that patient data privacy is protected, making drug development more efficient, secure, and compliant with global data privacy and security regulations[10]. For example, a prominent case currently being implemented in China is the federal learning framework of WeBank. By establishing a federal learning framework, WeBank can cross the risk control and financial credit assessment of different institutions without transferring the original data of customers. This approach not only protects the data privacy of financial clients, but also improves the predictive power of AI models, and most importantly, this approach ensures the privacy and security of all data.

As can be seen from several success stories above, building an enterprise's data consortium, coupled with incentives and possibly blockchain technology, can further simplify the process. These alliances will foster collaboration, ensure contributions are recognized and all participants benefit equitably, ultimately resulting in a more connected and efficient AI ecosystem. Webank's practice provides a model for other industries, demonstrating the practical potential of joint learning for data privacy protection.

3. Methodology

While Federated learning (FL) can provide a high level of security in terms of privacy protection, there are still some risks, such as reconstructing a single training model through model backward inference. One response is to inject noise and distort updates during the training of each node to hide the contribution of individual model nodes and limit the granularity of information shared between training nodes. However, existing research on privacy protection has focused on common machine learning benchmark datasets (such as MNIST) and stochastic gradient descent algorithms.

In this study, we implemented and evaluated a federal learning system for drug development data sharing. By experimenting with clinical trial data, we demonstrate the feasibility of medical data privacy protection technology in drug development.

Our key contributions include: (1) To our knowledge, the implementation and evaluation of the first privacy-protected federal learning system for drug development data analysis; (2) The use of joint average algorithm to deal with momentum optimization and unbalanced training nodes is compared; (3) The sparse vector technique (SVT) is empirically studied to obtain a strong differential privacy guarantee.

3.1. Federated learning framework

This paper uses the joint average algorithm to investigate a federated learning system based on a client-server architecture (shown on the left in Figure 5), where a central server maintains a global DNN model and coordinates local random gradient descent (SGD) updates on the client. This section describes the training process of the client model, the aggregation process of the server model, and the deployment of the privacy protection module on the client side.

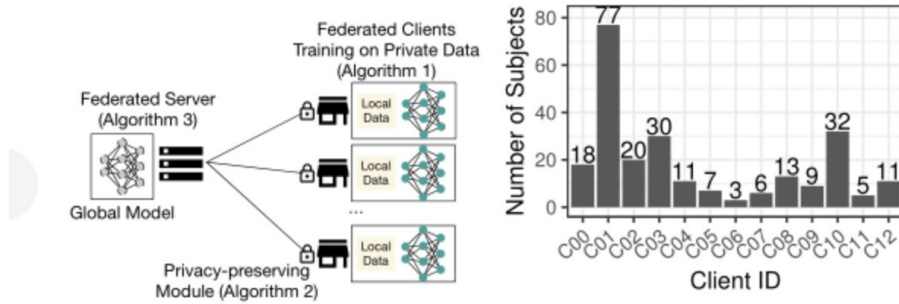


Figure 3: Left: illustration of the federated learning system; right: distribution of the training subjects ($N=242$) across the participating federated clients ($C=13$) studied in this paper.

3.2. Patient data model training process

We assume that each institution participating in federated learning has a fixed local data set and sufficient computing resources to conduct small-batch SGD updates. Clients share the same DNN structure and loss function. In round t of joint training, the local model is initialized by reading the global model parameter $w(t)$ from the server and updated to $w(l,t)$ through multiple SGD iterations. After a fixed number of local iterations, the model differences $\Delta w(t)$ are shared with the aggregation server.

In drug development, clinical trial data is often optimized using momentum-based SGD. Selective parameter updating: At the end of client training, the complete model may overfit and remember local training data, and sharing such a model may lead to data leakage. Therefore, the selective parameter

sharing method limits the amount of information shared by clients. Clients upload only a portion of $\Delta w(t)_k$, which is shared only if the parameter w_i is greater than the threshold $\tau(t)_k$. In addition, data privacy is further protected by clipping its values into a fixed range. The combination of clip gradient and selective parameter sharing can further enhance differential privacy through SVT.

3.3. Experimental result

Compared with centralized data set training, federated learning systems can also achieve better model performance without sharing customer data. In the drug development scenario, FL model training, despite a longer convergence time (about 600 rounds), still guarantees similar performance to a centralized dataset model. In addition, in the experiment, FL training time depends on the computing speed of the slowest client.

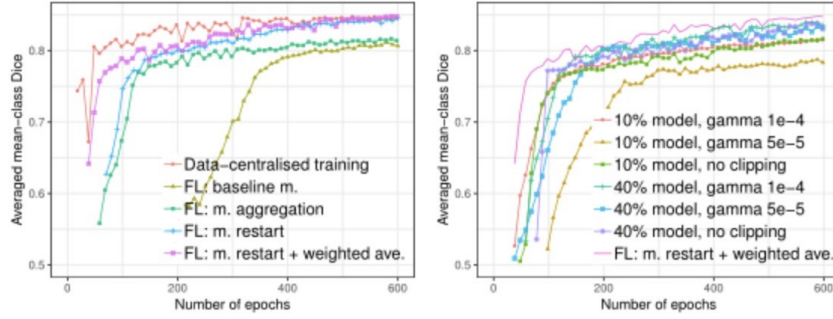


Figure 4: Comparison of segmentation performance on the test set with (left): FLvs. non-FL training, and(right): partial model sharing

4. Conclusion and Discussion

Through this study, we demonstrate the great potential of Federated learning (FL) in privacy-protected medical data sharing, especially in the field of drug development. Our experimental results show that despite significant differences in the data characteristics of different clients, the federated learning system is able to effectively train the model across multiple independent institutions without centralized sharing of sensitive patient data.

Momentum restart and weighted average: From the experimental results, the restart strategy for momentum variables significantly improves the convergence rate of the model, proving the necessity of restarting momentum in each round. This strategy avoids the interference of momentum variables between clients, thus ensuring the stability of the global model. By controlling for the share of parameters protected by DP, we found that sharing fewer model parameters performed better at the same privacy cost. This provides key implications for privacy protection in drug development - by optimizing the proportion of parameters shared, the best balance between privacy protection and model performance can be achieved.

Implications for drug development: Data security and privacy are important considerations in the drug development process, especially in the clinical trial phase. Federal learning technologies allow different pharmaceutical companies and research institutions to share clinical trial data without compromising patient privacy, accelerating the drug discovery and development process. For example, pharmaceutical companies can jointly develop more accurate drug response models while maintaining data localization, improving the efficiency of new drugs to market. Future research can explore adaptive model aggregation strategies to update personalized models according to the specific characteristics of clients, so as to improve the performance of global models. In addition, the further

optimization of differential privacy technology is also an important direction for future research, especially in large-scale drug development data, how to achieve stronger privacy protection without significantly affecting the model performance.

This study demonstrates the great potential of federal learning in the field of drug development. Through techniques such as momentum restart, local model sharing, and differential privacy, we successfully demonstrated how to achieve efficient model training while protecting data privacy. Future research could further optimize communication efficiency, model aggregation strategies, and privacy protection techniques to advance the application of federated learning in real-world drug development.

References

- [1] Bakas, S., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. *arXiv:1811.02629* (2018).
- [2] Hitaj, B., Ateniese, G., Perez-Cruz, F.: Deep models under the GAN: information leakage from collaborative deep learning. In: *SIGSAC*. pp. 603–618. *ACM* (2017)
- [3] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv:1412.6980* (2014)
- [4] Li, L., Fan, Y., Tse, M., & Lin, K. Y. (2020). A review of applications in federated learning. *Computers & Industrial Engineering*, 149, 106854.
- [5] Geyer, R.C., Klein, T., Nabi, M.: Differentially private federated learning: A client level perspective. *arXiv:1712.07557* (2017).
- [6] Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019, November). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM workshop on artificial intelligence and security* (pp. 1-11).
- [7] Xu, K., Zhou, H., Zheng, H., Zhu, M., & Xin, Q. (2024). Intelligent Classification and Personalized Recommendation of E-commerce Products Based on Machine Learning. *arXiv preprint arXiv:2403.19345*.
- [8] Xu, K., Zheng, H., Zhan, X., Zhou, S., & Niu, K. (2024). Evaluation and Optimization of Intelligent Recommendation System Performance with Cloud Resource Automation Compatibility.
- [9] Zheng, H., Xu, K., Zhou, H., Wang, Y., & Su, G. (2024). Medication Recommendation System Based on Natural Language Processing for Patient Emotion Analysis. *Academic Journal of Science and Technology*, 10(1), 62-68.
- [10] Zheng, H.; Wu, J.; Song, R.; Guo, L.; Xu, Z. Predicting Financial Enterprise Stocks and Economic Data Trends Using Machine Learning Time Series Analysis. *Applied and Computational Engineering* 2024, 87, 26–32.