

Enhancing Convergence in Federated Learning: A Contribution-Aware Asynchronous Approach

Changxin Xu^{1,*}, Yuxin Qiao², Zhanxin Zhou³, Fanghao Ni⁴, Jize Xiong⁵

¹ Computer Information Technology, Northern Arizona University, Flagstaff, USA

² Big Data and Business Intelligence, Universidad Internacional Isabel I de Castilla, Burgos, Spain

³ Computer Information Technology, Northern Arizona University, Flagstaff, USA

⁴ Software Engineering, Indian Institute of Technology, Guwahati, Assam, India

⁵ Computer Information Technology, Northern Arizona University, Flagstaff, USA

Abstract: Federated Learning (FL) is a distributed machine learning paradigm that allows clients to train models on their data while preserving their privacy. FL algorithms, such as Federated Averaging (FedAvg) and its variants, have been shown to converge well in many scenarios. However, these methods require clients to upload their local updates to the server in a synchronous manner, which can be slow and unreliable in realistic FL settings. To address this issue, researchers have developed asynchronous FL methods that allow clients to continue training on their local data using a stale global model. However, most of these methods simply aggregate all of the received updates without considering their relative contributions, which can slow down convergence. In this paper, we propose a contribution-aware asynchronous FL method that takes into account the staleness and statistical heterogeneity of the received updates. Our method dynamically adjusts the contribution of each update based on these factors, which can speed up convergence compared to existing methods.

Keywords: Asynchronous Federated Learning · Machine Learning · Distributed Computing · Privacy and Security.

1. Introduction

Over the past few years, Machine Learning has become a hot research topic. Federated Learning (FL) is a privacy-preserving distributed machine learning paradigm where clients utilize their local private data to collaboratively train a global model under the coordination of a centralized server. Most existing FL methods are developed under a synchronous communication setting, in which clients can upload their local models to the server simultaneously. It implicitly requires the assumption that all selected clients can finish the local training steps at roughly the same time. This assumption, however, hardly holds in many real-world systems. The devices typically have different computation resources and communication bandwidths, making the local updates received by the server asynchronous. The devices that finish the local training and upload to the server lately are defined as stragglers. A naive method to compromise the scenario is to request the faster clients waiting for the stragglers to participate in the aggregation. Although this method can somehow balance the contributions of each device, it largely decays the convergence speed of the whole FL system and causes computation resource waste on the faster devices.

To overcome the straggler problem in synchronous FL, several asynchronous FL methods have been developed to improve the convergence performance. In the fully asynchronous FL methods, the server performs a global model update once it receives any one local information without any waiting. Although this kind of fully asynchronous method helps increase system efficiency, it could decay the global model performance since the training bias is introduced during each global update. In addition, too frequent global aggregations cause a waste of server resources. Therefore, has proposed another line of asynchronous FL, in

which the server updates the global model once receiving the gradient from the first K clients' update while allowing the stragglers to continue their local training. Although this method achieves comparable performance with synchronous FL, it fails to measure the reasonable contributions of each local update during the aggregation step, as the server equally averages the received M local updates. The nature of different computation and communication abilities and statistical data heterogeneity cause the contributions of each received update different. In this work, we propose a systematic method to measure the contribution of each local update, aiming to address the following problems:

Problem 1: Statistical Heterogeneity. Statistical Heterogeneity (Non-IID) is considered as a challenging problem in federated learning. This problem is even accumulated in asynchronous federated learning scenarios. As devices communicate with the server under different frequencies, the global model would easily be biased to the faster devices. For each global model update, the weights of devices that contain the classes of data seldom studied should be higher than the devices with well-studied data.

Problem 2: Staleness. The heterogeneous computation and communication resources of each device create the problem of staleness. Without loss of the generality, assume for each client i which contributes to the t -th server update, the local update uploaded is trained with global model version t' and the staleness of this client i is defined as $\tau_i(t) = t - t'$. Intuitively, the gradients with low stale are relatively more reliable, while the gradients with high stale are more risky to borrow biased information. Thus, the existing asynchronous methods down-weight the contribution of gradient based on the staleness (e.g. $s(\tau_i(t)) = \frac{1}{(1+\tau_i(t))^{0.5}}$).

Although this kind of staleness-based weighted aggregation somehow reduces the risk of biased information, it could over-deteriorate the importance of high-stale

gradients from two aspects. On the one hand, measuring contributions which depend on the delayed time to upload, can neglect the global model differences. For a simple example of a two-class classification task, we consider that there are two clients, which client 1 has data with class A and client 2 has data with class B. If client 2 last communicates with the server when the server has already learned enough information from class A and the global model has been approaching stable on class A, the global model would change little, nevertheless the staleness of client 2's communication. In this case, downweighting the contribution of client 2's gradient is meaningless. On the other hand, reducing the weight of the high-stale gradient further increases the risk of statistical training bias, as the slow devices could contain data that has been hardly studied.

2. Related Work

2.1. Federated Learning

Federated Learning (FL) is a distributed machine learning framework that allows clients to collaboratively train a global model using their private data, under the coordination of a centralized server. Early works on FL, such as Federated Averaging (FedAvg), focused on reducing the communication overhead between the server and clients by using local stochastic gradient descent. Subsequent works have improved upon FedAvg by addressing the negative effects of statistical heterogeneity among clients. However, most of these works assume a synchronous communication setting, ignoring the system heterogeneity between clients. To address the issue of system heterogeneity, two types of methods have been proposed. One approach focuses on reducing the local computation workload of slower devices through partial model training in FL. In this paper, we focus on the other type of method, which allows faster and slower devices to communicate with the server asynchronously. This approach allows slower devices to continue training on their local data using a stale global model, without delaying the convergence of the overall FL system.

2.2. Asynchronous Federated Learning

Most existing works on asynchronous distributed learning, such as assume IID data distributions among the clients, which does not hold in the case of FL due to its inherent statistical heterogeneity. Other works in asynchronous federated learning, such as make additional assumptions, such as broadcasting model updates to all clients or assuming all clients have the same speed, to facilitate asynchronous FL.

One recent line of research on asynchronous FL updates the global model as soon as it receives any local information from clients, known as fully asynchronous FL. Although this kind of fully asynchronous FL helps to increase the system efficiency, it could enhance the cost of the server, as too frequent server aggregations are required. In addition, as the server aggregates each client's update, privacy could be a concern.

To solve the above problem, the most recent work studies another line of asynchronous FL, where the server updates the global model once receiving the gradient from the first K client's updates. However, these works equally aggregate the received gradient and ignore the different contributions among the gradients. In this work, we propose a novel FL method that takes into account the contributions of each update and adjusts them accordingly. Our method uses a

combination of weighting based on staleness and statistical heterogeneity to improve convergence performance.

3. Model and Preliminaries

We consider a FL system with one server and N clients. The clients work together to train a machine learning model by solving a distributed optimization problem as follows:

$$\min_x f(x) = \frac{1}{N} \sum_{i=1}^N f_i(x) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\zeta_i} [F_i(x, \zeta_i)], \quad (1)$$

where $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$ is a non-convex loss function for client i and F_i is the estimated loss function based on a mini-batch data sample ζ_i drawn from client i 's own dataset. Denote the dataset size of client i as D_i .

The classic fully asynchronous FL methods update the global model once the server receives any local updates from the client. Although this kind of method shows comparable convergence speed with synchronous settings, it raises privacy concerns. Recently, proposes a new line of asynchronous FL, in which the server has a buffer to store the received local updates. The server will only update the global model when the buffer contains K clients' local updates. (Note that the buffer size K is a hyper-parameter.) This kind of method is suitable to combine with the secured aggregation methods to prevent clients' privacy. Thus, in this paper, we choose to follow the main structure of.

Let Δ_t^i be the local update uploaded by client i to participate in the t round global aggregation. Δ_t^i is the cumulative gradients of M local training steps with initial model of global version $t - \tau_t^i$. In FedBuff, once the buffer collect K local updates, it will aggregate the new global model x_{t+1} as following:

$$x_{t+1} = x_x - \eta_g \frac{1}{K} \sum_{i=1}^K \Delta_t^i, \quad (2)$$

where η_g is the global learning rate at the server side. Although FedBuff performs well under moderate conditions, its performance could highly drop as it ignores the difference among the contributions of K local updates.

4. Methodology

The difference in contributions among K local updates comes from two aspects. On the first aspect, the system heterogeneity makes the required times to finish the M local training different. Typically the fast devices finish the local training sooner and upload timely information to the server, while the slow devices need more time to finish the local training and thus upload stale cumulative gradients to the server. On the other aspect, the statistical heterogeneity makes the different importance of clients' local updates to the current round model aggregation. In the following, we propose a novel method to measure the contribution of each client's local updates.

4.1. Staleness Effect

We here denote τ_t^i as the staleness between the model version in which local client i uses to start local training, and the model version in which client i participates in the t round global aggregation. Previous works down-weight stale updates using the function which is inversely propitiatory to

τ_i^t . (e.g., $s(\tau_i(t)) = \frac{1}{(1+\tau_i(t))^{0.5}}$). We measure the client i 's staleness degree S_i^t in t -th global aggregation as following:

$$S_i^t = \frac{\min_{j \in \mathcal{K}} \|x^t - x^{t-\tau_j^t}\|^2}{\|x^t - x^{t-\tau_i^t}\|^2}, \quad (3)$$

where \mathcal{K} is the set of local clients stored in the buffer to participate in the current round aggregation. Note that $S_i^t \in [0, 1]$. The client i that starts local training with the model version $x^{t-\tau_i^t}$ most similar to the current model version x^t has the highest $S_i^t = 1$.

4.2. Statistical Effect

Although the staleness effect is an important component in asynchronous FL scenarios, only applying staleness when considering the local clients' contributions could induce more problems. For example, consider a slow device that contains the classes of data which has been rarely learned by the server, simply applying staleness degree S_i^t could exacerbate the statistical heterogeneity. Thus, when we measure the contribution of clients' local updates, we should consider the clients' local statistical distributions. Unfortunately, privacy is a serious problem in the FL setting, which makes it impossible to directly report the local data distribution to the server. Moreover, as the statistical benefit of each client's local data to the server, keeps changing during each global model update, we require the 'fresh' information which is calculated based on the current global model version to determine the statistical effect. Contributes that in client selection, selecting the client with higher local loss could enhance the convergence speed. We propose to utilize the batch loss to estimate the local training loss. Typically, the higher training loss represents a worse learning of the current global model on local data. Thus, we define the statistical effect of client i in t -th global aggregation as:

$$P_i^t = \mathcal{N}_i \frac{1}{|\zeta_i|} F_i(x^t, \zeta_i) \quad (4)$$

where ζ_i is the mini-batch data sample randomly drawn from client i 's local dataset.

Then we can update the global model as follows:

$$x_{t+1} = x_t - \eta_g \frac{1}{K} \sum_{i=1}^K P_i^t \Delta_t^i / S_i^t, \quad (5)$$

where η_g is the global learning rate at the server side.

5. Experiment

We conduct experiments for the image classification task on Fashion-MNIST [34] for non-IID data distributions. There are 30 clients and each client has 1500 instances for training. In each FL round, all clients participate. We use LeNet as the backbone model.




Fig. 1. Performance Comparison

The result shown in Fig. 1 validates that our proposed method outperforms the baseline method by a large margin.

6. Conclusion

In conclusion, this paper introduces a novel contribution-aware asynchronous Federated Learning (FL) method that represents a significant advancement in distributed machine learning. By addressing the limitations of traditional synchronous FL methods, which often suffer from slow and unreliable communication in realistic settings, our approach offers a more efficient and practical solution. Unlike existing asynchronous FL methods that aggregate updates indiscriminately, our method considers the staleness and statistical heterogeneity of each update, dynamically adjusting their contributions to the global model. This nuanced approach ensures faster convergence and enhances the overall performance of FL systems. Our contribution-aware method not only mitigates the challenges posed by asynchronous updates but also paves the way for more robust and scalable FL implementations. Future research could further refine this approach, exploring additional factors that influence the effectiveness of updates and extending our method to a wider range of applications. Through continued innovation in FL, we can better harness the power of distributed machine learning while upholding the essential principles of privacy and efficiency.

References

- [1] Ye, J., Kang, H., Wang, H., Shen, C., Jahannia, B., Heidari, E., ... & Dalir, H. (2023, September). Demultiplexing OAM beams via Fourier optical convolutional neural network. In *Laser Beam Shaping XXIII* (Vol. 12667, pp. 16-33). SPIE.
- [2] Ye, J., Solyanik, M., Hu, Z., Dalir, H., Nouri, B. M., & Sorger, V. J. (2023, March). Free-space optical multiplexed orbital angular momentum beam identification system using Fourier optical convolutional layer based on 4f system. In *Complex Light and Optical Forces XVII* (Vol. 12436, pp. 70-80). SPIE.
- [3] Ye, J., Kang, H., Wang, H., Altaleb, S., Heidari, E., Asadizanjani, N., ... & Dalir, H. (2023, October). OAM beams multiplexing and classification under atmospheric turbulence via Fourier convolutional neural network. In *Frontiers in Optics* (pp. JTU4A-73). Optica Publishing Group.
- [4] Ye, J., Kang, H., Wang, H., Altaleb, S., Heidari, E., Asadizanjani, N., ... & Dalir, H. (2023, November). Multiplexed OAM beams classification via Fourier optical convolutional neural network. In *2023 IEEE Photonics Conference (IPC)* (pp. 1-2). IEEE.
- [5] Kang, H., Ye, J., Wang, H., Dalir, H., & Sorger, V. J. (2023, November). Optical Freespace Michelson Interferometric

- Reconfigurable Full Complex Convolution Module. In 2023 IEEE Photonics Conference (IPC) (pp. 1-2). IEEE.
- [6] Kang, H., Solyanik-Gorgone, M., Ye, J., Nouri, B. M., Wang, H., Dalir, H., & Sorger, V. J. (2023, November). Secure Optical Hashing for Information Compression in a Convolutional Neural Network. In 2023 IEEE Photonics Conference (IPC) (pp. 1-2). IEEE..
- [7] Wang, L., Bian, J., & Xu, J. (2023). Federated Learning with Instance-Dependent Noisy Labels. arXiv preprint arXiv:2312.10324.
- [8] Bian, J., Ren, S., & Xu, J. (2023). CAFE: Carbon-Aware Federated Learning in Geographically Distributed Data Centers. arXiv preprint arXiv:2311.03615.
- [9] Bian, J., & Xu, J. (2023, October). Client Clustering for Energy-Efficient Clustered Federated Learning in Wireless Networks. In Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing (pp. 718-723).
- [10] Bian, J., Shen, C., & Xu, J. (2023). Joint Client Assignment and UAV Route Planning for Indirect-Communication Federated Learning. arXiv preprint arXiv:2304.10744.
- [11] Liu, T., Xu, C., Qiao, Y., Jiang, C., & Yu, J. (2024). Particle Filter SLAM for Vehicle Localization. arXiv preprint arXiv:2402.07429.
- [12] Liu, T., Xu, C., Qiao, Y., Jiang, C., & Chen, W. (2024). News Recommendation with Attention Mechanism. arXiv preprint arXiv:2402.07422.
- [13] Guo, C., Zhao, Y., Liu, T., & Yang, C. (2023). THE ROLE OF MACHINE LEARNING IN ENHANCING COMPUTER VISION PROCESSING. In АКТУАЛЬНЫЕ ВОПРОСЫ СОВРЕМЕННЫХ НАУЧНЫХ ИССЛЕДОВАНИЙ (pp. 23-26).
- [14] Popokh, L., Su, J., Nair, S., & Olinick, E. (2021, September). IllumiCore: Optimization Modeling and Implementation for Efficient VNF Placement. In 2021 International Conference on Software, Telecommunications and Computer Networks (SoftCOM) (pp. 1-7). IEEE.
- [15] Su, J., Nair, S., & Popokh, L. (2023, February). EdgeGym: A Reinforcement Learning Environment for Constraint-Aware NFV Resource Allocation. In 2023 IEEE 2nd International Conference on AI in Cybersecurity (ICAIC) (pp. 1-7). IEEE.
- [16] Su, J., Nair, S., & Popokh, L. (2022, November). Optimal Resource Allocation in SDN/NFV-Enabled Networks via Deep Reinforcement Learning. In 2022 IEEE Ninth International Conference on Communications and Networking (ComNet) (pp. 1-7). IEEE.
- [17] Xu, J., & Sen, S. (2023). Ensemble Variance Reduction Methods for Stochastic Mixed-Integer Programming and their Application to the Stochastic Facility Location Problem. *INFORMS Journal on Computing*.
- [18] Xu, J., & Sen, S. (2023). Compromise policy for multi-stage stochastic linear programming: Variance and bias reduction. *Computers & Operations Research*, 153, 106132.
- [19] Xu, J., & Sen, S. (2021). Decision Intelligence for Nationwide Ventilator Allocation During the COVID-19 Pandemic. *SN Computer Science*, 2(6), 423.
- [20] Deng, Y., Kesselman, C., Sen, S., & Xu, J. (2019, December). Computational operations research exchange (core): A cyber-infrastructure for analytics. In 2019 Winter Simulation Conference (WSC) (pp. 3447-3456). IEEE.
- [21] Liu, S., Wu, K., Jiang, C., Huang, B., & Ma, D. (2023). Financial Time-Series Forecasting: Towards Synergizing Performance And Interpretability Within a Hybrid Machine Learning Approach. arXiv preprint arXiv:2401.00534.
- [22] Xiong, J., Feng, M., Wang, X., Jiang, C., Zhang, N., & Zhao, Z. (2024). Decoding Sentiments: Enhancing COVID-19 Tweet Analysis through BERT-RCNN Fusion. *Journal of Theory and Practice of Engineering Science*, 4(01), 86-93.
- [23] Xu, C., Yu, J., Chen, W., & Xiong, J. (2024, January). DEEP LEARNING IN PHOTOVOLTAIC POWER GENERATION FORECASTING: CNN-LSTM HYBRID NEURAL NETWORK EXPLORATION AND RESEARCH. In The 3rd International scientific and practical conference "Technologies in education in schools and universities"(January 23-26, 2024) Athens, Greece. International Science Group. 2024. 363 p. (p. 295).
- [24] Zang, H. (2024). Precision Calibration of Industrial 3D Scanners: An AI-Enhanced Approach for Improved Measurement Accuracy. *Global Academic Frontiers*, 2(1), 27-37.
- [25] Zang, H., & Dong, X. (2024). Optimizing Soil Health Management in Smart Agriculture: Deep Learning Algorithms for Nutrient Analysis and Fertilizer Recommendation with Precision Agriculture Systems. *Journal of Industrial Engineering and Applied Science*, 2(1), 1-7.
- [26] Qiao, Y., Jin, J., Ni, F., Yu, J., & Chen, W. (2023). APPLICATION OF MACHINE LEARNING IN FINANCIAL RISK EARLY WARNING AND REGIONAL PREVENTION AND CONTROL: A SYSTEMATIC ANALYSIS BASED ON SHAP. *WORLD TRENDS, REALITIES AND ACCOMPANYING PROBLEMS OF DEVELOPMENT*, 331.
- [27] Qiao, Y., Ni, F., Xia, T., Chen, W., & Xiong, J. (2024, January). AUTOMATIC RECOGNITION OF STATIC PHENOMENA IN RETOUCHEDED IMAGES: A NOVEL APPROACH. In The 1st International scientific and practical conference "Advanced technologies for the implementation of new ideas"(January 09-12, 2024) Brussels, Belgium. International Science Group. 2024. 349 p. (p. 287).
- [28] Ni, F., Zang, H., & Qiao, Y. (2024, January). SMARTFIX: LEVERAGING MACHINE LEARNING FOR PROACTIVE EQUIPMENT MAINTENANCE IN INDUSTRY 4.0. In The 2nd International scientific and practical conference "Innovations in education: prospects and challenges of today"(January 16-19, 2024) Sofia, Bulgaria. International Science Group. 2024. 389 p. (p. 313).
- [29] Xiaobin, L., Changlin, Y., & Wanting, W. (2023). CARASSISTNET: DESIGN AND IMPLEMENTATION OF A DRIVING ASSISTANCE SYSTEM USING COMPUTER VISION. *Ответственный редактор*, 28.
- [30] Hu, W., Liu, X., & Xie, Z. (2022). Ore image segmentation application based on deep learning and game theory. In *World science: problems and innovations* (pp. 71-76).