Federated Learning Based on Deep Reinforcement Learning

Subjects: Computer Science, Artificial Intelligence Contributor: Xutao Meng , Yong Li , Jianchao Lu , Xianglin Ren

Federated learning (FL) is a distributed machine learning paradigm that enables a large number of clients to collaboratively train models without sharing data.

federated learning deep reinforcement learning client selection

1. Introduction

The application of deep learning technology in the Internet of Things (IoT) is very common, with uses in smart healthcare, smart transportation, and smart cities ^[1]. However, the massive amounts of data in IoT impose limitations on traditional centralized machine learning in terms of network resources, data privacy, etc. The proposal of federated learning (FL) provides an effective solution to deep learning problems involving data privacy issues. Clients can collaborate with other clients in training a global model without the need to share their local data [2]. FL has been successfully applied in many domains ^{[3][4][5][6][7]}. However, the presence of data heterogeneity among clients adversely affects the model convergence and training accuracy of FL.

In real-world scenarios, the local datasets among different clients exhibit heterogeneity, indicating that their local data distribution differs from the global data distribution within the entire federated learning (FL) system. Several studies have demonstrated that the heterogeneity of data among clients significantly impacts the effectiveness of FL methods, leading to a substantial reduction in model accuracy ^{[8][9]}. Specifically, heterogeneity among clients can lead to inconsistent convergence targets for their local training. Aggregating local models with biased convergence targets will naturally result in a global model with biased convergence targets as well. Therefore, the divergence of the global model obtained from non-IID datasets as opposed to IID datasets continues to grow, which may lead to slower convergence and poorer learning performance ^[10]. Effectively mitigating the adverse effects of data heterogeneity on FL system models remains one of the central challenges in current federated learning optimization.

Some researchers consider only a single category of non-IID environments and do not provide stable performance improvements in different categories of non-IID environments [11][12][13][14][15][16]. Furthermore, The authors in [17] restrain the local model update and mitigate the degree of "client drift" by introducing a variety of proximal terms. Their approach of introducing proximal terms is effective, but also inherently limits the potential for local model convergence while incurring considerable communication, computation, and storage overhead on the clients, which are intolerable for real-world distributed systems. Moreover, most previous work [11][12][13][14][15][16][17][18][19]

assumes that all clients in the FL system can participate in each round of iteration. The number of clients is usually smaller when all clients participate, while in practice, the number of clients is typically larger. The scenario with a multitude of participants, where all clients participate in each FL round, is not feasible due to differences in communication or computational power, among other factors. In common approaches, the server typically employs a random selection policy to select participants and uses a model aggregation algorithm to update the weights of the global model. The random selection of clients for participation in the model aggregation process can increase the bias of the global model and exacerbate the negative effects of data heterogeneity. Therefore, it is crucial to design an optimal client selection strategy that is robust for FL.

Numerous studies focus on devising client selection strategies to alleviate the issue of data heterogeneity in FL. Some authors measure the degree of local dataset skews by utilizing the discrepancy between local and global model parameters for the development of client selection strategies ^{[20][21][22]}. These methods either rely on a global shared dataset or cause a huge waste of resources. The training loss generated during local training naturally reflects the degree of skew in the local data distribution and training progress between different clients. Other than that, The calculation and upload of the loss value will not generate new computational or storage burdens. Some studies use that biased selection based on client-side local training loss values and achieve good results ^{[23][24][25]}. They believe that favoring clients with higher local loss values can accelerate the convergence of FL in heterogeneous environments. Intuitively, in the early stages of FL, clients with high loss values will help the global model converge faster. However, choosing clients with high loss values may negatively impact accuracy improvement when the global model is close to convergence.

Deep reinforcement learning (DRL) excels at handling optimal decisions in complex dynamic environments, where the agent repeatedly observes the environment, performs actions to maximize its goals, and receives rewards from the environment. Constructing an agent for the server in FL, the agent adaptively selects clients with high or low loss values to participate in the global model aggregation process by designing a suitable reward function, thus alleviating the problem that client selection strategies are difficult to formulate in dynamic environments.

2. Data-Based Approaches

Several studies have attempted to alleviate the non-IID issue among clients. Zhao et al. ^[26] improved training on non-IID data by constructing a small, globally shared, uniformly distributed data subset for all clients. Similarly, Seo et al. ^[27] mitigated the quality degradation problem in FL via data sharing, using an auction approach to effectively reduce the cost, while satisfying system requirements for maximizing model quality and resource efficiency. In ^[28], the authors assume that a small segment of clients are willing to share their datasets, and the server collects data from these clients in a centralized manner to aid in updating the global model. Although such data-sharing-based methods have obtained significant performance improvements, they go against the original intention of FL and pose a threat to privacy. And in the absence of the client's original data, the server cannot obtain the global data distribution information and use it to build a globally shared IID dataset.

3. Algorithm-Based Approaches

Another research aspect focuses on addressing the negative impact of heterogeneous data by designing algorithms to enhance the local training phase or improve the global aggregation process. In [11], the authors introduce a new algorithm called SCAFFOLD. The algorithm uses control variables to correct for local updates, preventing "client drift", and leverages the similarity in client data to accelerate the convergence of FL. Li et al. [12] balances the optimization differences between global and local objectives using a regularization term. In addition, the authors [13] introduced a normalized averaging algorithm called FedNove. This algorithm normalizes local updates by the number of local training iterations per client. It ensures rapid error convergence while maintaining objective consistency. The authors of ^[14] propose the FedRS method, which constrains the updates of missing category weights during local training via a classification layer in a neural network. MOON ^[15] is proposed as model-contrastive federated learning. It introduces a contrastive loss for the clients, utilizing the representations of the global model and historical local models for learning, to correct the local model updates of each client. Similarly, the authors of ^[16] proposed FedProc, a prototypical contrastive federated learning approach. The authors design a global prototypical contrastive loss for local network training and use prototypes as global knowledge to correct local training for each client. The authors of [18] demonstrate a contribution-dependent weighting design, named FedAdp. It calculates the association between the client's local goals and the global goal of the overall FL system based on the gradient information during the training process, assigning different weights to each participating client. Zhang et al. ^[19] address the challenge of direct model aggregation by transferring knowledge from the local model to the global model through data-free distillation. Long et al. ^[29] propose FedCD, which removes classifier bias from non-IID data by introducing hierarchical prototype comparison learning, global information distillation, and other methods to understand the class distribution of clients.

4. System-Based Approaches

In addition, several studies have attempted to design client selection policies for servers. In ^[20], the authors determine the level of IID data among clients by analyzing differences in local model weights. They assign a higher probability of selection to clients with lower degrees of non-IID, ensuring their more frequent participation in FL training. But the assumption of accessible IID public data is challenging to meet in the real world. Wu et al. ^[21] use the inner product of the local model gradient and the global model gradient as a measure to determine the subset of clients participating in model aggregation, ensuring that clients contributing more to reducing the global loss have a higher probability of being selected. Some studies have designed a client selection strategy by considering the local training loss values. Goetz et al. ^[25] evaluate the contribution of different client data in each FL round according to the local loss value, calculate the corresponding evaluation score, and select an optimized subset of clients according to the evaluation value. Cho et al. ^[23] theoretically demonstrate that favoring client selection with larger local loss values can improve the convergence rate compared to random client selection. Other studies employ reinforcement learning to select clients for servers. Chen et al. ^[30] use an UCB approach to heuristically select participating clients during each round of optimization, utilizing the cosine distance weights (CDW) of the historical global model and the current local model to measure the client's contribution and assign rewards.

Moreover, the author of ^[22] proposed an experience-driven control framework that uses a deep reinforcement learning algorithm to intelligently select clients in each round of federated learning (FL) by reducing the dimensionality of clients' local model weights and using them as states to enhance the performance of the global model. Xiao et al. ^[31] proposed a client selection strategy based on clustering and bi-level sampling. Firstly, a subset of candidate clients is constructed using MD sampling, and then a WPCS mechanism is proposed to collect the weighted per-label mean class scores of the clients to perform clustering and select the final client.

References

- 1. Nguyen, D.C.; Ding, M.; Pathirana, P.N.; Seneviratne, A.; Li, J.; Poor, H.V. Federated learning for internet of things: A comprehensive survey. IEEE Commun. Surv. Tutorials 2021, 23, 1622–1658.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B.A.Y. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, Lauderdale, FL, USA, 20–22 April 2017; pp. 1273–1282.
- Yu, Z.; Hu, J.; Min, G.; Wang, Z.; Miao, W.; Li, S. Privacy-preserving federated deep learning for cooperative hierarchical caching in fog computing. IEEE Internet Things J. 2021, 9, 22246–22255.
- 4. Venkataramanan, V.; Kaza, S.; Annaswamy, A.M. Der forecast using privacy preserving federated learning. IEEE Internet Things J. 2022, 10, 2046–2055.
- Zhang, C.; Liu, X.; Zheng, X.; Li, R.; Liu, H. Fenghuolun: A federated learning based edge computing platform for cyber-physical systems. In Proceedings of the 2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops), Austin, TX, USA, 23–27 March 2020; pp. 1–4.
- 6. He, Y.; Huang, K.; Zhang, G.; Yu, F.R.; Chen, J.; Li, J. Bift: A blockchain-based federated learning system for connected and autonomous vehicles. IEEE Internet Things J. 2021, 9, 12311–12322.
- 7. Wang, K.; Deng, N.; Li, X. An efficient content popularity prediction of privacy preserving based on federated learning and wasserstein gan. IEEE Internet Things J. 2022, 10, 3786–3798.
- 8. Hsu, T.-M.H.; Qi, H.; Brown, M. Measuring the effects of non-identical data distribution for federated visual classification. arXiv 2019, arXiv:1909.06335.
- 9. Li, X.; Huang, K.; Yang, W.; Wang, S.; Zhang, Z. On the convergence of fedavg on non-iid data. arXiv 2019, arXiv:1907.02189.
- 10. Zhu, H.; Xu, J.; Liu, S.; Jin, Y. Federated learning on non-iid data: A survey. Neurocomputing 2021, 465, 371–390.

- Karimireddy, S.P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; Suresh, A.T. Scaffold: Stochastic controlled averaging for federated learning. In Proceedings of the International Conference on Machine Learning, Vienna, Austria, 12–18 July 2020; pp. 5132–5143.
- 12. Li, T.; Sahu, A.K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; Smith, V. Federated optimization in heterogeneous networks. Proc. Mach. Learn. Syst. 2020, 2, 429–450.
- 13. Wang, J.; Liu, Q.; Liang, H.; Joshi, G.; Poor, H.V. Tackling the objective inconsistency problem in heterogeneous federated optimization. Adv. Neural Inf. Process. Syst. 2020, 33, 7611–7623.
- Li, X.-C.; Zhan, D.-C. Fedrs: Federated learning with restricted softmax for label distribution noniid data. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, Singapore, 14–18 August 2021; pp. 995–1005.
- Li, Q.; He, B.; Song, D. Model-contrastive federated learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 10713–10722.
- 16. Mu, X.; Shen, Y.; Cheng, K.; Geng, X.; Fu, J.; Zhang, T.; Zhang, Z. Fedproc: Prototypical contrastive federated learning on non-iid data. Future Gener. Comput. Syst. 2023, 143, 93–104.
- Li, Q.; Diao, Y.; Chen, Q.; He, B. Federated learning on non-iid data silos: An experimental study. In Proceedings of the 2022 IEEE 38th International Conference on Data Engineering (ICDE), Kuala Lumpur, Malaysia, 9–12 May 2022; pp. 965–978.
- 18. Wu, H.; Wang, P. Fast-convergent federated learning with adaptive weighting. IEEE Trans. Cogn. Commun. Netw. 2021, 7, 1078–1088.
- Zhang, L.; Shen, L.; Ding, L.; Tao, D.; Duan, L.-Y. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 10174–10183.
- 20. Zhang, W.; Wang, X.; Zhou, P.; Wu, W.; Zhang, X. Client selection for federated learning with noniid data in mobile edge computing. IEEE Access 2021, 9, 24462–24474.
- 21. Wu, H.; Wang, P. Node selection toward faster convergence for federated learning on non-iid data. IEEE Trans. Netw. Sci. Eng. 2022, 9, 3099–3111.
- 22. Wang, H.; Kaplan, Z.; Niu, D.; Li, B. Optimizing federated learning on non-iid data with reinforcement learning. In Proceedings of the IEEE INFOCOM 2020-IEEE Conference on Computer Communications, Toronto, ON, Canada, 6–9 July 2020; pp. 1698–1707.
- Cho, Y.J.; Wang, J.; Joshi, G. Towards understanding biased client selection in federated learning. In Proceedings of the International Conference on Artificial Intelligence and Statistics, Valencia, Spain, 28–30 March 2022; pp. 10351–10375.

- 24. Cho, Y.J.; Wang, J.; Joshi, G. Client selection in federated learning: Convergence analysis and power-of-choice selection strategies. arXiv 2020, arXiv:2010.01243.
- 25. Goetz, J.; Malik, K.; Bui, D.; Moon, S.; Liu, H.; Kumar, A. Active federated learning. arXiv 2019, arXiv:1909.12641.
- 26. Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; Chandra, V. Federated learning with non-iid data. arXiv 2018, arXiv:1806.00582.
- 27. Seo, E.; Niyato, D.; Elmroth, E. Resource-efficient federated learning with non-iid data: An auction theoretic approach. IEEE Internet Things J. 2022, 9, 25506–25524.
- Yoshida, N.; Nishio, T.; Morikura, M.; Yamamoto, K.; Yonetani, R. Hybrid-fl for wireless networks: Cooperative learning mechanism using non-iid data. In Proceedings of the ICC 2020–2020 IEEE International Conference On Communications (ICC), Dublin, Ireland, 7–11 June 2020; pp. 1–7.
- 29. Long, Y.; Xue, Z.; Chu, L.; Zhang, T.; Wu, J.; Zang, Y.; Du, J. Fedcd: A classifier debiased federated learning framework for non-iid data. In Proceedings of the 31st ACM International Conference on Multimedia, Ottawa, ON, Canada, 29 October–3 November 2023; pp. 8994–9002.
- Chen, W.; Zhou, X. Ucbfed: Using reinforcement learning method to tackle the federated optimization problem. In Proceedings of the Distributed Applications and Interoperable Systems: 21st IFIP WG 6.1 International Conference, DAIS 2021, Held as Part of the 16th International Federated Conference on Distributed Computing Techniques, DisCoTec 2021, Valletta, Malta, 14– 18 June 2021; Springer: Berlin/Heidelberg, Germany, 2021; pp. 99–105.
- Xiao, D.; Zhan, C.; Li, J.; Wu, W. Bi-level sampling: Improved clients selection in heterogeneous settings for federated learning. In Proceedings of the 2023 IEEE International Performance, Computing and Communications Conference (IPCCC), Anaheim, CA, USA, 17–19 November 2023; pp. 430–435.

Retrieved from https://encyclopedia.pub/entry/history/show/117518