

Comparative Analysis of Personalized Federated Learning Optimization Algorithms for Image Classification

Can Liao

*School of Computer Science and Technology, Huazhong University of Science and Technology,
Wuhan, China
U202215414@hust.edu.cn*

Abstract: As data privacy becomes more vital and data heterogeneity prevails in image classification, personalized federated learning optimization algorithms have come to the fore as an essential solution. These algorithms enable multiple clients to train personalized models while maintaining the privacy of their data, thus enhancing the performance of image classification. This study is targeted at conducting a thorough comparison among various personalized federated learning optimization algorithms when it comes to image classification. The proposed method follows a comparative study framework, where a global model is initialized and made available to multiple clients. Each client trains a personalized model using specific algorithms that incorporate both local data and the global model. The server then aggregates model updates according to the respective rules until convergence, with accuracy serving as the primary performance metric. Experiments were performed using the Canadian Institute for Advanced Research (CIFAR)-10 dataset, with the outcomes revealing varying test accuracies for algorithms as the number of clients changes. The findings demonstrate that each algorithm handles data heterogeneity and client numbers differently, showcasing their respective strengths and weaknesses in terms of accuracy, overfitting prevention, and adaptability to local data. These insights provide a solid foundation for selecting appropriate algorithms in practical scenarios.

Keywords: Personalized Federated Learning, Data Privacy, Image Classification, Federated Learning.

1. Introduction

Recently, as data privacy gains more significance, federated learning has come to the fore as a feasible machine-learning methodology. This method enables numerous clients to work together to train a global model while keeping their original private data secure, safeguarding user privacy in the process [1]. However, traditional federated learning, aiming to create a one-fits-all solution, endeavors to train a global model that is supposed to be applicable to all clients, often overlooking the diversity of client-side data. In the context of image classification, data heterogeneity is very common, as different clients may have different distributions of image data. optimization algorithms for personalized federated learning have been put forward to tackle this problem. Their objective is to train customized models for each client, thus enhancing the performance of image classification [2].

In the field of personalized federated learning, various optimization algorithms have been proposed. Here is a brief introduction to some of them. Personalized Federated Averaging (Per-

FedAvg) adjusts the global model aggregation method. After the standard FedAvg process, it conducts multiple rounds of local fine-tuning. While it is straightforward to implement, this approach might cause overfitting issues [3]. Federated Adaptive Model Personalization (FedAMP) introduces personalized parameters during weighted aggregation on the server side to make the global model adapt to the characteristics of each client, achieving better personalization with less local data [4]. Personalized Federated Learning with Moreau Envelopes (pFedMe) uses the Moreau Envelopes technology to balance personalization and global information by constraining the local model update, being compatible with non-convex optimization problems and having good theoretical convergence [5]. Federated Proximal (FedProx) incorporates an extra L2 regularization term during the local update phase. This is done to restrict the disparity between the local and global models, effectively managing heterogeneous data and minimizing the model drift problem [6]. Federated Representation (FedRep) splits the model into a shared feature extraction layer and a personalized classification layer, only updating the personalized part locally, which is suitable for complex tasks and can reduce communication overhead [7]. Adaptive Personalized Federated Learning (APFL) allows each client to train both the local model and the global model simultaneously, dynamically adjusting the personalization degree through adjustable parameters [8]. Personalized Federated Learning using Hypernetworks (pFedHN) uses a HyperNetwork to generate personalized model parameters, being suitable for resource-constrained devices and reducing communication costs [9]. Personalized Federated Reinforcement Learning (PerFedRL) combines federated learning and meta-learning to make the reinforcement learning model adapt to different client task goals and environments, being suitable for reinforcement learning tasks [10]. Federated Expectation-Maximization (FedEM) trains multiple personalized models using the EM algorithm and uses a weighted voting strategy for prediction during inference, having good theoretical convergence [11].

The primary objective of this study is to compare various personalized federated learning optimization algorithms for image classification. To achieve this, state-of-the-art personalized federated learning algorithms are selected and implemented. The performance of these algorithms is then evaluated and compared based on accuracy, loss, and convergence speed, using the same image dataset. Additionally, the study explores the impact of different hyperparameters on algorithm performance to identify optimal parameter settings. The experimental results reveal that certain algorithms outperform others in specific scenarios, such as achieving higher accuracy or faster convergence. This study offers valuable insights for selecting the most suitable personalized federated learning optimization algorithms for practical applications, thereby enhancing the efficiency and accuracy of image classification while preserving privacy.

2. Methodology

2.1. Dataset description and preprocessing

This research uses the Canadian Institute for Advanced Research (CIFAR)-10 dataset [2]. It consists of 60,000 color images categorized into 10 classes, with 6,000 images per class. This dataset is commonly employed in the image classification domain for model training and evaluation. Before training, the data undergoes pre-processing prior to training. Specifically, pixel values are normalized to the range of [0, 1], and the dataset is partitioned into a training set, a validation set, and a test set according to a specific ratio.

2.2. Proposed approach

Conducting a comprehensive comparison of various personalized federated learning optimization algorithms for image classification is what this research primarily aims to do. The research methodology includes implementing several typical personalized federated learning algorithms. Each

of these algorithms is assessed using a set of standard performance indicators. As illustrated in Figure 1, First of all, the initialization of the global model takes place, followed by its distribution to all the participating clients. After that, each client proceeds to train its personalized model, which is based on its local data combined with the global model, and applies the specific personalized federated learning algorithm in the process. After the clients finish their training, the server aggregates the models or parameter updates from these clients. This aggregation is carried out in accordance with the algorithm's pre-set rules. This iterative process is repeated until convergence criteria are satisfied. Finally, the performance of the trained models is assessed on a test set, using evaluation metrics such as accuracy, loss, and convergence speed.




Figure 1: The pipeline of the model (picture credit: original)

2.2.1. Per-FedAvg

Per-FedAvg is a personalized federated learning algorithm that modifies the traditional global model aggregation strategy. After the standard FedAvg process, it conducts local fine-tuning for a predefined number of rounds. This approach is based on the idea that after the initial global training, local data can be further exploited to enhance the model's performance on each client. The standard FedAvg training rounds are first executed by clients when the initial global model is sent to them. This global training helps the model capture general features across all clients. Subsequently, clients enter the local fine-tuning stage. During this stage, the model is adjusted based on the local data distribution. For example, if a client has a unique set of images related to a particular class in the CIFAR-10 dataset, the local fine-tuning can enhance the model's capacity for recognizing these specific images. However, this also brings a risk. If the local data amount is relatively large, the model may overfit to the local data. This overfitting occurs because the model becomes too specialized to the local patterns and may not generalize well to other clients' data. Despite this risk, Per-FedAvg's simplicity makes it a practical choice in many scenarios, especially when the local data is not extremely large, and quick adaptation to local data is required.

2.2.2. FedProx

FedProx incorporates an additional L2 regularization term during the local update process. The core objective of including this term is to reduce the disparity between the local model and the global model. Mathematically, the local objective function of FedProx can be formulated as equation (1).

$$L_{local}^k(\theta) = \mathcal{L}(f(x_{local}^k; \theta), y_{local}^k) + \frac{\mu}{2} \|\theta - \theta_{global}\|^2 \quad (1)$$

As a constraint, this term serves to restrict the disparity between the local and global models. When clients utilize their own data to train local models, they minimize the L2-regularized local objective function. By incorporating this regularization term into the loss function, FedProx effectively mitigates the model drift problem among clients. In a scenario where clients have heterogeneous data, this helps in maintaining the stability of the overall federated learning process. For instance, if some clients have data with a different distribution from the global average, the L2 regularization term guarantees that the local models won't diverge too much from the global model, consequently improving the generalization ability of the model for image classification tasks. However, in some cases, the L2 regularization may be too restrictive, potentially preventing the model from fully capturing the unique characteristics of the local data.

2.2.3. FedRep

FedRep employs a distinctive method of dividing the model into a shared feature extraction layer and a personalized classification layer. For image classification, the shared feature extraction layer extracts common features from images, and these features can be shared among all clients. This layer captures common characteristics that are relevant for classifying images in the dataset, such as edges, textures, and basic shapes. On the other hand, the personalized classification layer is updated locally based on each client's data. During the training process, clients focus on updating this personalized part while keeping the shared feature extraction layer fixed during certain stages. This design has several advantages. Firstly, it serves to reduce the communication overhead that exists between clients and the server, which is beneficial for improving the overall efficiency of the system. Since only the personalized classification layer's updates need to be transmitted, less data is sent, which is especially beneficial in scenarios with limited network bandwidth. Secondly, it enhances the personalized effect. By adapting the classification layer to the local data, the model can better classify images according to the specific characteristics of each client's data. However, as the number of clients changes and the data distribution becomes more diverse, FedRep may face challenges in effectively balancing the global information from the shared layer and the local information from the personalized layer.

2.2.4. pFedHN

pFedHN leverages a Hypernetwork to generate personalized model parameters, which renders it especially appropriate for resource-constrained devices, like those in edge computing scenarios.

The Hypernetwork is trained to generate appropriate parameters for each client's local model by considering both the global information and the client's own data characteristics. As shown in Figure 2, this process allows clients to obtain models that are well-adapted to their specific data distributions without consuming excessive communication resources. In the CIFAR-10 dataset experiment, the Hypernetwork analyzes the global trends in the data as well as the distinctive features of local data. It then generates personalized parameters that can optimize the performance of the local model on the client's data. For example, if a client has a higher proportion of a certain class of images, the Hypernetwork can generate parameters that are more sensitive to the features of that class. This personalized parameter generation not only improves the model's accuracy on the local data but also

reduces the communication cost, as the generated parameters are tailored to the client's needs and can be used directly without extensive additional communication. However, training the Hypernetwork requires careful parameter adjustment. If the parameters are not set correctly, the Hypernetwork may not fully utilize its potential for generating highly personalized models, leading to sub-optimal performance.




Figure 2: Illustration of pFedHN mechanism (picture credit: original)

2.3. Implementation details

The system is developed using the Python programming language and the PyTorch library. For data augmentation, such methods as random flipping, rotation as well as cropping are adopted to expand the diversity of the training data. Hyperparameters including the learning rate, the number of local training epochs, the batch size, and the communication rounds between clients and the server are meticulously adjusted via a series of preliminary experiments. This is done to guarantee that each personalized federated learning algorithm achieves optimal performance in the image classification task.

3. Results and discussions

3.1. Results analysis

As presented in Table 1, the test accuracy of four personalized federated learning algorithms on the CIFAR10 dataset across different numbers of clients (10, 50, 100) is observed.

When the number of clients is 10, Per-FedAvg attains an accuracy of 76.65. FedProx shows a higher accuracy of 87.27, FedRep reaches 87.69, and pFedHN performs best among them with 90.83. As the number of clients raises to 50, Per-FedAvg's accuracy rises to 83.03, FedProx slightly decreases to 83.39, FedRep is at 86.09, and pFedHN maintains a high level at 88.38. When the number of clients further increases to 100, Per-FedAvg's accuracy is 80.19, FedProx goes up to 89.99, FedRep drops to 85.23, and pFedHN stands at 87.97.

The diverse performance of these algorithms can be ascribed to their different approaches in dealing with data heterogeneity. An analysis shows that the local fine-tuning in Per-FedAvg enables it to adapt to local data. However, when the local data volume is large, overfitting may occur, accounting for its performance fluctuations. FedProx's L2 regularization term effectively curbs model

drift as the number of clients escalates, leading to more stable performance improvements. FedRep's split-model design balances global and local information, yet it faces challenges in fully adapting when client numbers grow. pFedHN, by leveraging the Hypernetwork, efficiently generates personalized parameters, enabling it to maintain high accuracy across different client numbers. This analysis is crucial for practical applications, as it guides the selection of algorithms based on the scale of clients and data heterogeneity.

Table 1: Test accuracy of different algorithms on CIFAR10 dataset

Algorithm	CIFAR10 (10 clients)	CIFAR10 (50 clients)	CIFAR10 (100 clients)
Per-FedAvg	76.65	83.03	80.19
FedProx	87.27	83.39	89.99
FedRep	87.69	86.09	85.23
pFedHN	90.83	88.38	87.97

3.2. Discussion

Per-FedAvg is characterized by its ease of implementation. Its local fine-tuning mechanism allows the model to adapt to local data patterns. However, it has a tendency to overfit when the local data volume is substantial, restricting its effectiveness in scenarios with extensive and diverse local datasets. FedProx benefits from the L2 regularization term that mitigates model drift for better handling of heterogeneous data, though it may be overly restrictive in some cases and impede capturing local data characteristics. FedRep's split model design reduces communication overhead and enhances personalization, suitable for complex image classification but struggles to balance global and local information with varying client-side data distributions. pFedHN uses a Hypernetwork for personalized parameter generation, which suits resource-constrained devices and cuts communication costs, though training it demands careful parameter adjustment or else may not fully utilize its potential for personalized model generation.

In terms of future research directions, there's a need to develop more sophisticated hyperparameter-tuning mechanisms for these algorithms to adapt better to real-world data dynamics. Exploring algorithm integration, like combining FedProx's regularization with pFedHN's parameter generation, could yield more robust solutions. Also, investigating their performance in non-iid data with complex distributions is essential to expand the applicability of personalized federated learning in diverse data scenarios.

4. Conclusion

This study investigates the performance of various personalized federated learning optimization algorithms in image classification, leveraging the CIFAR10 dataset. The main objective is to surmount the challenges of data heterogeneity in federated learning and identify algorithms that can enhance both the efficiency and accuracy of image classification while ensuring data privacy. The proposed method utilizes a comparative study framework, beginning with the initialization of a global model, and then distributed to client devices. Each client trains a personalized model based on local data and the global model. The model updates are aggregated at the server in accordance with the specific rules of each algorithm. This iterative procedure lasts until convergence. Accuracy serves as the primary metric for performance evaluation. Through comprehensive experiments, the study evaluates the four personalized federated learning algorithms. The results reveal distinct performance characteristics for each algorithm when handling data heterogeneity and varying client numbers. The pros and cons of each algorithm are apparent in terms of accuracy, overfitting prevention, and adaptability to local data distributions.

Future research will explore the development of automated hyperparameter-tuning mechanisms to optimize these algorithms further. Additionally, the combination of different algorithms and their performance in more complex non-iid data settings represent promising directions for future investigation.

References

- [1] Liu, J., Huang, J., Zhou, Y., Li, X., Ji, S., Xiong, H., & Dou, D. (2022). From distributed machine learning to federated learning: a survey. *Knowledge and Information Systems*, 64(4), 885 - 917.
- [2] Tan, A. Z., Yu, H., Cui, L., & Yang, Q. (2023). Towards Personalized Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12), 9587 - 9603.
- [3] Beaussart, M., Grimberg, F., Hartley, M. A., & Jaggi, M. (2021). Waffle: Weighted averaging for personalized federated learning. *arXiv preprint arXiv:2110.06978*.
- [4] Pei, J., Liu, W., Li, J., Wang, L., & Liu, C. (2024). A review of federated learning methods in heterogeneous scenarios. *IEEE Transactions on Consumer Electronics*.
- [5] T Dinh, C., Tran, N., & Nguyen, J. (2020). Personalized federated learning with moreau envelopes. *Advances in neural information processing systems*, 33, 21394-21405.
- [6] Mishchenko, K., Khaled, A., & Richtárik, P. (2022). Proximal and federated random reshuffling. In *International Conference on Machine Learning*, 15718-15749.
- [7] Jing, C., Huang, Y., Zhuang, Y., Sun, L., Xiao, Z., Huang, Y., & Ding, X. (2023). Exploring personalization via federated representation Learning on non-IID data. *Neural Networks*, 163, 354-366.
- [8] Zhang, J., Hua, Y., Wang, H., Song, T., Xue, Z., Ma, R., & Guan, H. (2023). Fedala: Adaptive local aggregation for personalized federated learning. In *Proceedings of the AAAI conference on artificial intelligence*, 37(9), 11237-11244.
- [9] Shamsian, A., Navon, A., Fetaya, E., & Chechik, G. (2021). Personalized Federated Learning using Hypernetworks. *Proceedings of the International Conference on Machine Learning*, 9489-9502.
- [10] Roth, K., Akrou, R., Gruenewalder, S., & Nowozin, S. (2020). PerFedRL: A Personalized Federated Reinforcement Learning Framework. *arXiv preprint*, 6440.
- [11] Zhao, Y., Li, M., Lai, L., & Liang, Y. (2020). Federated EM: A Federated Expectation-Maximization Framework for Privacy-Preserving Learning. *arXiv preprint*, 6440.