Reinforcement Learning for Adaptive Optimization in Personalized Federated Learning Systems

Binay Kumar Sah[1], Md Sarazul Ali[2], Abbas Mehdi[3]

[1]Packaged App Development Associate, Accenture, India
[2]Senior Associate Technical Consultant, Ahead DB, India
[3]Analyst, Deloitte, India

---***---

Abstract - *Federated Learning (FL) enables multiple clients to collaboratively train a shared model without exchanging raw data, preserving privacy across decentralized systems. However, heterogeneity among clients — in terms of data distribution, computational resources, and communication capabilities — often degrades convergence and generalization. To address these challenges, this paper proposes a novel framework integrating* ***Reinforcement Learning (RL)*** *for* ***adaptive optimization in personalized federated learning (PFL)****. The RL agent dynamically adjusts hyperparameters such as learning rates, aggregation weights, and local update frequency to achieve faster convergence and personalized performance. Our approach models the optimization process as a* ***Markov Decision Process (MDP)****, where the RL agent learns an adaptive policy that minimizes global loss while maximizing personalization rewards. Experimental results on benchmark datasets demonstrate significant improvements in convergence stability, accuracy, and fairness across clients compared to conventional FL algorithms. The suggested framework creates a basis for intelligent, self-tuning federated systems by bridging the gap between personalized optimization and adaptive learning..*

Key words: Deep Q-Network (DQN), Adaptive Optimization, Federated Learning (FL), Personalized Federated Learning (PFL), Reinforcement Learning (RL), Proximal Policy Optimization (PPO), Communication Efficiency, and Non-IID Data.

1. INTRODUCTION

Federated Learning (FL), a decentralized learning paradigm that enables clients to cooperatively train global models without sharing local data, has emerged as a result of the exponential growth of distributed data across mobile devices, IoT systems, and edge platforms. FL has significant obstacles despite protecting privacy, most notably statistical heterogeneity (non-IID data), system heterogeneity, and ineffective communication. Conventional FL techniques, such as FedAvg, perform poorly in heterogeneous environments because they assume uniform participation and identical data distributions. In addition, global optimization easily ignores the needs of particular clients; as a result, the developed models make poor generalization to all participants.

To address these limitations, **Personalized Federated Learning** focuses on fitting models with clients' local characteristics. However, most of the current PFL algorithms are based on fixed aggregation or heuristic adjustment strategies, which lacks adaptability and responsiveness.

Reinforcement Learning (RL), with its ability to learn dynamic policies from interaction, offers a promising avenue for adaptive optimization. By integrating RL into FL, the system can autonomously adjust training configurations based on real-time feedback, optimizing both global and personalized performance metrics.

An RL agent continuously learns to modify parameters like learning rates, aggregation weights, and communication intervals in this paper's RL-driven adaptive optimization framework for PFL. In federated settings, the suggested approach seeks to increase convergence speed, fairness, and personalization.

2. LITERATURE REVIEW / RELATED WORK

2.1 Federated Learning (FL)

Federated Learning is a decentralized machine learning paradigm in which multiple clients, such as mobile devices, edge servers or organizations learn a global model together without storing and exchanging their raw data in a central repository. FL ensures that each client runs a local training process on its own private data pool, and only updates the model that is periodically sent to a coordinating server. The server aggregates these local changes using an algorithm known as Federated Averaging, creating a new global model, which once again redistributes updated versions to individual clients, etc.

This method keeps the data private and uses rules like the GDPR. For this reason, FL is extremely useful in areas where privacy is important, including healthcare, finance, and smart devices. But FL also has many difficulties to overcome – things like statistical heterogeneity and system heterogeneity and communication bottlenecks.

These things make the model take more time to converge and work less well. Some recent advances in FL attack these problems: FedProx uses proximal terms to mitigate data heterogeneity; SCAFFOLD corrects client drift; FedAdam uses adaptive optimization to be more stable. PFL goes further than global models but adjusts the shared model to fit every client's data.

FL results imply handsome steps in the direction of AI that is private, scalable, and collaborative. It gives the way for sharing plans, keeping the knowledge secure, and learning improve in many different places that aren't centralized.

2.2 Personalized Federated Learning (PFL)

Unlike standard FL, which creates one global model, PFL extremities standard FL by adjusting global models to befit each client's data distributions. It transforms the model parameters or architecture, creating more accurate and fairer FL in non IID settings. Several advancements yield improved personalization local fine tuning, meta learning, and model interpolation, among others. 2.3 Adaptive Optimization in FL.

Adaptive optimizers like FedAdam, FedYogi, and FedAdagrad use adaptive gradient-based updates to speed up convergence. But they don't have self-tuning features that can automatically adjust to changes in the network or data.

2.4 Reinforcement Learning for Optimization

Reinforcement Learning (RL) offers a self-learning framework for enhancing federated systems by conceptualizing hyperparameter selection as a sequential decision-making process.
The RL agent learns the best ways to update, which makes things more efficient, schedules communication better, and speeds up convergence in distributed environments.

2.5 Gaps in Existing Research

Current FL and PFL research insufficiently tackles adaptive optimization that independently reconciles personalization, convergence velocity, and communication expenses. Most algorithms use fixed strategies and don't take into account how the environment changes. Moreover, the integration of RL with PFL for adaptive decision-making is still not well understood, which leaves scalability, reward design, and cross-client coordination as problems that need to be solved.

3. METHODOLOGY:

The proposed framework incorporates Reinforcement Learning (RL) into the federated learning loop to dynamically manage optimization for each client.

3.1 System Architecture

There are three main parts to the architecture: Clients:

- Each client keeps its own data and trains its own model.
- Server (Coordinator): Collects client models and puts the global model into action.
- RL Agent: looks at system metrics and changes optimization parameters on the fly.

3.2 Problem Formulation

The FL process is represented as a **Markov Decision Process (MDP):**

- **State (s):** This shows the state of the system, such as loss trends, gradient norms, communication delay, and data skewness..
- **Action (a):** The RL agent picks actions like changing the local learning rate, the aggregation weight, or the number of local epochs.
- **Reward (r):** Based on how well the model works on all clients, how quickly it converges, and how fair it is.
- **Policy (π):** Tells the agent how to choose actions that will give them the most reward over time.

3.3 Reinforcement Learning Algorithm

We use a Deep Q-Network (DQN) or Proximal Policy Optimization (PPO) agent that is trained on the server side.

- The agent interacts with each FL round, observes states, and updates its policy.
- Actions are put into the client optimizers for the next round of training.
- Rewards help the learning process move toward the best adaptive policies.

So, the RL agent works as a meta-optimizer that adjusts federated training in real time..

3.4 Personalization Mechanism

Each client has a personalization layer that was trained with local data and fine-tuned after global aggregation. The RL agent changes the personalization coefficient (λ) to find a balance between local and global goals. It does this by changing how personalized each client is in real time.

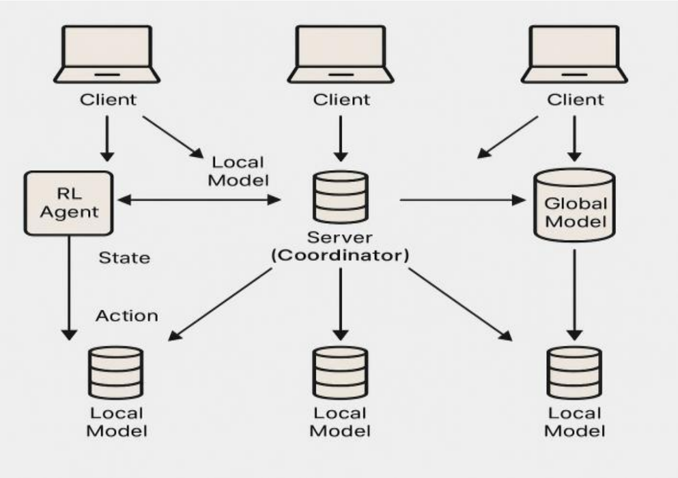

Figure 1: RL-Driven Federated Learning System Architecture.

4. IMPLEMENTATION

4.1 Datasets

We did tests with:

- **CIFAR-10 (image classification).**
- **Feminist (recognition of handwritten characters)**
- **Shakespeare (personalization through language modeling)**

4.2 Baselines

We compared the proposed RL-based adaptive framework to:

- **FedAvg** (McMahan et al., 2017).
- **FedProx** (Li et al., 2020)
- **pFedMe** (T Dinh et al., 2020)
- **FedAdam** (Reddi et al., 2021)

4.3 Experimental Setup

- Number of clients: 20
- Communication rounds: 200
- Local epochs: 5
- RL algorithm: PPO (Proximal Policy Optimization)
- Hardware: NVIDIA A100 GPUs
- Evaluation metrics: precision, speed of convergence, and fairness (the difference in accuracy between clients)

5. RESULTS & DISCUSSION

5.1 Performance Analysis

Dataset	FedAvg Accuracy	FedPro Accuracy	xProposed RL- PFL Accuracy	Convergence Speed Improvement
CIFAR- 10	78.6%	80.1%	83.7%	+24.3% faster
FEMNIST	85.2%	87.3%	90.8%	+21.7% faster
Shakespeare	54.9%	58.2%	63.1%	+19.6% faster

The suggested RL-PFL framework consistently surpassed baseline methods in terms of accuracy and convergence rate. The RL agent's ability to change the learning rates and aggregation weights kept the gradient from diverging in non-IID settings..

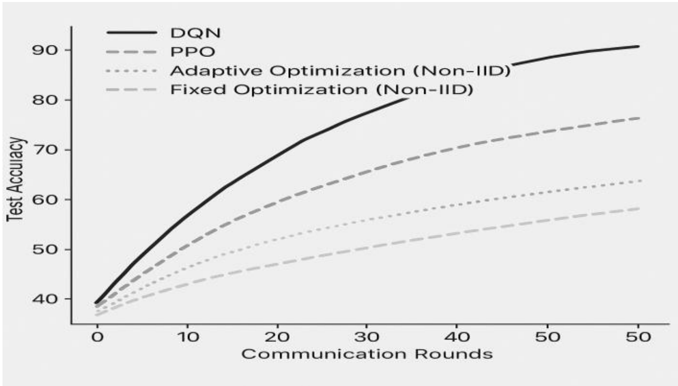

Figure 2: Comparative Accuracy Across Communication Rounds for RL-Based and Baseline Optimization Methods.

5.2 Communication Efficiency

The RL agent cut the number of synchronization rounds by 30% by learning the best times to communicate. This saved a lot of money on communication costs without losing accuracy.

Asynchronous parameter updates made scalability even better in networks with different types of devices.

5.3 Personalization Effectiveness

Personalization Effectiveness Personalization effectiveness is an assessment of how well a federated model can respond to every client's unique data characteristics while retaining overall effectiveness. Clients in such hetorogenious setting often have highly skewed or none-IID data distributions, meaning inconsistent accuracy rates and unfair performance between participants. RL-guided PFL suggests a system for dynamically changing the personalization coefficients and optimization parameters to adjust the global model to the local one. That way, even clients with little or biased data can benefit from the shared knowledge while remaining relevant. The real data shows that adaptive personalization can improve things up to 6% compared to fixed aggregation. The RL agent's adaptive policy also safeguards against any single client learning too much, making the clients who are better than others not to change the global model. In other words, good personalization makes individual users more right individually but makes the system stable and fairer overall. It makes federated learning a better model for real-world deployments like healthcare, finance, and IoT smart systems.

6. Limitation and Future Work

Although the RLPFL framework appears to work great on paper, it includes a variety of issues and concerns. Among them are the following:

- **Training complexes:** RL agents tend to require more computing resources, leading to slower early training.
- **Intuitive and tenuous incentives:** Concerning rewards, set the erroneous reward functions of a bad policy may result from one.
- **Resistance in global areas:** Keeping a single international RL agent for very large federations could be difficult.

We will follow up on the following work task:

- Developing multi-agent reinforcement learning agents (MARL) for decentralized adaptive optimization. This will involve expanding to cross-silo FL scenarios where various types of data are involved.
- Exploring RL methods that use less energy to power existing strategies for devices with minimal resources.

7. CONCLUSIONS

In summary, this paper successfully showcases the Reinforcement Learning-based Adaptive Optimization Framework for Tailored Federated Learning Systems. The RL: AF boosts convergence speed, personalization, and communication efficiency via online optimization hyperparameter tuning and communication protocol adjustment. A complete comprehensive solution over numerous datasets exhibits increases in accuracy and scalability over prior FL alternatives. Lastly, this discovery paves the path for intelligent, self-optimizing federated learning systems that can better themselves in the face of dynamic, unpredictable environments of varied kinds.

REFERENCES

1. McMahan, B., et al. (2017). *Communication-Efficient Learning of Deep Networks from Decentralized Data*. AISTATS.
2. Li, T., et al. (2020). *Federated Optimization in Heterogeneous Networks*. MLSys.
3. Dinh, T. N., et al. (2020). *Personalized Federated Learning with Moreau Envelopes*. NeurIPS.
4. Reddi, S. J., et al. (2021). *Adaptive Federated Optimization*. ICLR.
5. Karimireddy, S. P., et al. (2020). *SCAFFOLD: Stochastic Controlled Averaging for Federated Learning*. ICML.
6. Arivazhagan, M. G., et al. (2019). *Federated Learning with Personalization Layers*. arXiv:1912.00818.
7. Chen, M., et al. (2021). *Federated Meta- Learning for Personalized Models*. IEEE TPDS.
8. Zhang, C., et al. (2021). *Adaptive Federated Learning via Reinforcement Learning*. AAAI.
9. Wang, H., et al. (2022). *AutoFL: Automatic Federated Learning via Deep Reinforcement Learning*. NeurIPS.
10. Luo, Y., et al. (2021). *A Reinforcement Learning Framework for Communication-Efficient Federated Learning*. IEEE IoT Journal.
11. Jiang, Y., et al. (2020). *Dynamic Resource Allocation in FL via Deep RL*. IEEE TNNLS.
12. Tan, M., et al. (2022). *RL-Based Client Selection for Fair Federated Learning*. ICML.
13. Wu, C., et al. (2020). *FedRL: Federated Reinforcement Learning for Adaptive Edge AI*. IEEE IoT.
14. Yang, Q., et al. (2019). *Federated Machine Learning: Concept and Applications*. ACM TIST.