The Multi-Arm Bandit Problem in the Federated Context

A Technical Report submitted to the Department of Computer Science

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Aaron parson

Spring, 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Briana Morrison, Department of Computer Science

The Multi-Arm Bandit Problem in a Federated Context

CS 4991 Capstone Report, 2024

Aaron Parson
The University of Virginia
School of Engineering and Applied Science
Charlottesville, Virginia USA
asp6qw@virginia.edu

ABSTRACT

Data sharing, privacy concerns, and infrastructure costs have become notable problems in machine learning and limit data diversity. Federated Learning (FL) operates as a distributed system of devices to help solve these machine-learning problems. The specific problem addressed in FL is the Multi-Armed Bandit problem (MAB), a reinforcement learning problem where the objective is to identify a slot machine with the highest true reward while minimizing regret out of a certain number of other slot machines. Using the FedML platform, the MAB problem was placed into a Federated Context allowing for increased model performance and data sharing. Future work could explore the significance of the increase in data sharing and data diversity of the network and how to maintain security in the network from adversarial bandit devices.

1. INTRODUCTION

Last year, I conducted Machine Learning Research on Federated Learning. Federated Learning (FL) is an emerging technology that addresses various limitations in machine learning, including data sharing, privacy concerns, and infrastructure costs, while enhancing data diversity. Removing these limits on machine learning models can allow for better performance, thereby improving the potential for what is possible in AI and analytics fields. While this technology can aid in the battle in maintaining consumer privacy, it is one steppingstone in a much larger solution. In this paper, I want to explore my project in Federated Learning and how it can be used to help improve consumer privacy.

2. BACKGROUND

Federated Learning creates a distributed learning environment on multiple devices working to solve machine learning problems (Zhang & Bai et al 2021). This network includes multiple client devices, each creating their own machine learning models on data local to the device. The client devices will send their model updates to a global server in which these updates are aggregated into the global model. After the aggregation of the model updates is complete, the global server will then send the aggregated model back to the client devices. The client devices will then integrate the global model into their own local model and then continue training their new model on local data. This process is repeated for a desired number of interactions creating a distributed machine learning network. This technology is privacy preserving because the data never leaves the device that it was collected on. Therefore, there is less risk in data sharing. However, there are other risks and issues including

privacy protection, data sufficiency, and statistical heterogeneity. Still, FL has a multitude of applications, one of which being the integration with the Multi-Arm Bandit Problem (MAB). The MAB problem is a reinforcement learning problem in which the goal is to estimate the slot machine with the highest average reward from a varying number of slot machines. The notable concepts are the reward which is the amount received from playing a certain number of slots and the regret which is the difference between the reward received and the true highest reward. In this problem, the reward is to be maximized and the regret is to be minimized. Placing the MAB problem in a FL network allows for each device playing in the MAB problem to share information between each other. This benefits each player's ability to estimate the slot machine with the true highest average reward. This therefore creates a distributed recommending system in which privacy can be better protected (Shi & Shen 2021).

3. PROJECT DESIGN

The initial phase of the project involved conducting an extensive review of the current literature on federated machine learning (FML). Through examination of academic papers, industry reports, and relevant research articles, built a comprehensive understanding of FML methodologies, challenges, and advancements. This process allowed for the identification key trends, emerging techniques, and areas of interest within the FML landscape, laying the foundation for our research.

The evaluation of the Intel federated machine learning (FML) package was conducted with attention to various criteria essential for its applicability in research endeavors. The assessment encompassed a comprehensive review of the package's features and capabilities, guided by specific parameters tailored to project requirements.

Commencing with scrutiny of the package's support for custom federated learning (FL) models, the aim was to understand its flexibility in accommodating tailored models suited to research objectives. Furthermore, exploration extended to the availability of pre-implemented state-of-the-art FL models within the package, recognizing the potential for leveraging existing advancements in the field to expedite the research process.

Additionally, delving into the package's provisions for custom hyperparameter tuning was crucial for optimizing model performance according to specific research requirements. Addressing concerns related to data heterogeneity and privacy, assessment of the package's specifications for handling separate local data and mechanisms for client-side data manipulation ensued, ensuring compliance with privacy regulations, and maintaining data integrity.

Analysis of the package's support for both asynchronous and synchronous communication protocols was pivotal for accommodating diverse network conditions and communication requirements in federated learning setups. Compatibility with multiple processor architectures was also scrutinized to ensure scalability and deployment flexibility across various hardware environments.

Moreover, rating the ease of modification for both hardware and software components on a scale of 1-5 provided insight into the package's adaptability to evolving research needs. Deployment flexibility examination aimed

to streamline the research workflow and facilitate experimentation.

Lastly, verification of the package's open-source nature underscored its transparency, community collaboration, and long-term sustainability. Through meticulous evaluation of these criteria, valuable insights into the Intel FML package's suitability for research purposes were gained, ultimately informing the decision-making process and shaping the trajectory of federated machine learning investigations.

Based on our evaluation results and project requirements, the decision to utilize the FedML platform for the research was made. The FedML platform offered a comprehensive set of tools and functionalities tailored specifically to federated learning research, making it an ideal choice for our project goals. With the platform selected, focus transitioned towards implementing the distributed linear Upper Confidence Bound (UCB) algorithm within the FedML framework. This phase began with in-depth research into the theoretical underpinnings of the Multi-arm bandit problem and its pertinence to federated learning. Phd students had already implemented simulation code for the distributed multi-arm bandit problem. This algorithm was then integrated into the FedML platform.

4. RESULTS

Following the successful implementation of the distributed linear Upper Confidence Bound (UCB) algorithm within the FedML framework, the simulation was able to seamlessly integrate with the existing codebase. However, it is imperative to conduct further validation to measure how the simulation differs and aligns with the actual implementation within the code. Additionally, scalability testing is crucial to ascertain the algorithm's performance under varying conditions and dataset sizes. Rigorous testing and validation procedures are needed to evaluate not only the algorithm's performance but also its scalability and efficacy in real-world federated learning applications. Through meticulous experimentation and thorough analysis, the algorithm's suitability and effectiveness within the federated learning framework can be comprehensively assessed. These findings will lay a solid foundation for future advancements and wider applications in the field of federated machine learning.

5. RELATED WORKS

FL can be a possible solution to the lag in legislation protecting customer privacy. Because of legislation such as HIPPA, data accessibility in the health care industry is strict. Having increased accessibility to data can support medical research on a higher level (Kaissis & Makowski et al, 2020). Machine learning models, which thrive on diverse data sets, are unable to improve accuracy. Therefore, other research works to use federated learning as a solution to share medical information. It centers around mitigating privacy issues using federated learning and creating a secure and private way to share medical information. Specifically, the research involves maintaining anonymity and pseudonymity while limiting the risk of reidentification. More research works on attacking Federated Learning networks (Sun & Dong at el, 2021). One specific attack that relates to federated learning is a poisoning attack. This attack tries to skew model performance by altering the data of one or

more nodes in a federated machine learning network. From understanding how these attacks effect the network, algorithms can be created to mitigate the effects of the attacks.

6. CONCLUSION

Implementing MAB into FedML allowed for the creation of a distributed Reinforcement learning network. While fine-tuning is necessary, this project has laid a solid foundation for future FML research to be built from especially those dealing with reinforcement learning.

7. FUTURE WORK

Future work in this project would be fine-tuning the machine learning network to understand it accuracy and what hyperparameters are best for each device in the network to learn. Working to implement asynchronous learning is another future work in which this will help eliminate the need for the global server to wait for each device's models before aggregation. Other algorithms can be implemented in FedML allowing for testing of these algorithms in federated context.

8. ACKNOWLEDGEMENTS

Honging Wang, Ethan Blaser, Hongyan Wu, Matthew Whelan, Karan Singh, Xidong Wu

References

Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, *2*(6), 305–311. https://doi.org/10.1038/s42256-020-0186-1

Shi, C., & Shen, C. (2021). Federated Multi-Armed Bandits. *Proceedings of the AAAI Conference on Artificial Intelligence*, *35*(11), 9603–9611. https://doi.org/10.1609/aaai.v35i11.17156

Sun, G., Cong, Y., Dong, J., Wang, Q., Lyu, L., & Liu, J. (2021). Data Poisoning Attacks on Federated Machine Learning. *IEEE Internet of Things Journal*, 1–1. https://doi.org/10.1109/jiot.2021.3128646

Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems*, *216*, 106775. https://doi.org/10.1016/j.knosys.2021.106775