

# A Multi-player Game for Studying Federated Learning Incentive Schemes

Kang Loon Ng<sup>1</sup>, Zichen Chen<sup>2</sup>, Zelei Liu<sup>1</sup>, Han Yu<sup>1,2\*</sup>, Yang Liu<sup>3†</sup> and Qiang Yang<sup>3,4</sup>

<sup>1</sup>School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore

<sup>2</sup>Joint NTU-WeBank Research Centre on FinTech, NTU, Singapore

<sup>3</sup>Department of AI, WeBank, Shenzhen, China

<sup>4</sup>Department of Computer Science and Engineering, Hong Kong University of Science and Technology  
han.yu@ntu.edu.sg, yangliu@webank.com

## Abstract

Federated Learning (FL) enables participants to “share” their sensitive local data in a privacy preserving manner and collaboratively build machine learning models. In order to sustain long-term participation by high quality data owners (especially if they are businesses), FL systems need to provide suitable incentives. To design an effective incentive scheme, it is important to understand how FL participants respond under such schemes. This paper proposes *FedGame*, a multi-player game to study how FL participants make action selection decisions under different incentive schemes. It allows human players to role-play under various conditions. The decision-making processes can be analyzed and visualized to inform FL incentive mechanism design in the future.

## 1 Introduction

Artificial intelligence (AI) has enjoyed rapid development benefiting from availability of big data. Traditionally, data are stored in a centralized entity for training of machine learning models. However, centralized training may intrude user privacy as specified under the General Data Protection Regulation (GDPR) [Yang *et al.*, 2019b]. Federated Learning (FL) has been proposed as an alternative paradigm for building AI models in a distributed and privacy-preserving manner [McMahan *et al.*, 2016; Yang *et al.*, 2019a; Kairouz *et al.*, 2019]. It enables multiple participants to collaboratively train AI models without exposing potentially sensitive local data, thereby, improving users’ trust in the AI model [Pan *et al.*, 2009; Shen *et al.*, 2011; Yu *et al.*, 2018].

Participating in FL incurs costs, which can be significant for business participants [Yu *et al.*, 2020] under Horizontal Federated Learning (HFL), in which participants have significant overlap in the feature space but little overlap in the sample space [Yang *et al.*, 2019b]. These costs can arise from communication, technical, compliance, risk of market share erosion and free-riding problems (i.e. participants may only join FL training with low-quality data) [Yang *et al.*, 2019b].

\*Contact Author

†Contact Author

To sustain the long-term viability of FL ecosystems, effective incentive mechanisms are needed. To this end, the research community needs to understand how FL participants behave under given incentive schemes.

In this paper, we bridge this gap with *FedGame* - a multi-player game [Yu *et al.*, 2017] which aims to study how FL participants act under different incentive schemes through crowdsourcing [Pan *et al.*, 2016]. It supports multiple FL payoff-sharing schemes (currently including Linear, Equal, Individual, Labour Union and Shapley, with the possibility to extend to others) [Yang *et al.*, 2017; Gollapudi *et al.*, 2017; Jia *et al.*, 2019]. Through *FedGame*, researchers can analyze human players’ behaviours to improve FL incentive schemes.

## 2 System Architecture

Figure 1 illustrates the *FedGame* system architecture. In the game, a number of AI players and Federations are created to simulate the FL environment. A human player plays the role of a business (with an arbitrary amount of data and resources allocated to him at the beginning of the game) joining the federation. His business is assumed to be from the same market sector as the AI players’, which means contributing his data to the federation might result in an FL model that helps himself as well as his competitors [Yu *et al.*, 2020].

Key information such as resource quantity, data quality,

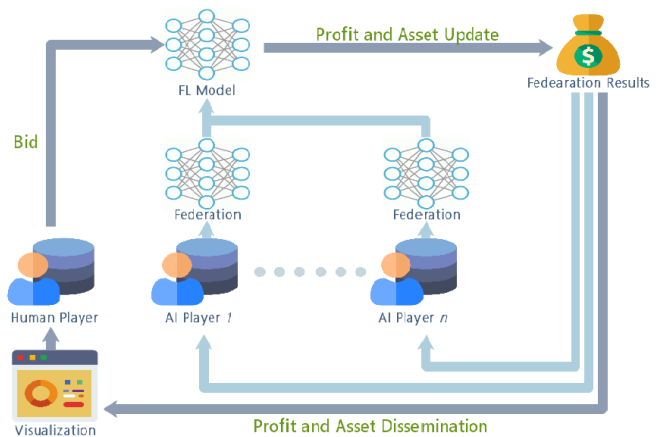


Figure 1: *FedGame* system architecture.

data quantity and payment are involved in decision-making. Multi-agent AI players [Yu *et al.*, 2010; Yu *et al.*, 2011; Wu *et al.*, 2013] follow existing approaches to determine how much data they want to contribute during FL model training [Yu *et al.*, 2020]. The human players decide how to allocate their resources that they want to contribute to the training of FL model based on individual free will. The Federations will receive payoffs from the virtual marketplace based on the market share their FL models occupy. Participants will be rewarded with a portion of their Federation’s payoff according to the incentive scheme adopted by Federation during a game session. The players’ in-game behaviour data are recorded.

### 3 Game Design

Each game instance ends after a fixed number of turns have passed. The ultimate goal for a player is to receive as much payoff as possible at the end of a game instance. In order to incentivize participants to contribute high-quality data to FL model training and truthfully report private cost types, the design of the game firstly focuses on illustrating the FL environment from the perspective of business enterprises. A player can decide to join, leave or remain in a Federation at any point in the game. The game system provides functions for game designers to modify existing incentive schemes or add new incentive schemes by creating new levels in the game. Each time when a player enters the game, he/she will be randomly assigned to a starting characteristic in terms of the amount and quality of the local data. This will be done through the randomization of allocated variables to players. With a different initialization at the start of every new game, players will not be lulled into following the same decision-making pattern but instead, be forced to adapt their behaviours.

Each Federation will be initialized with a fixed amount of credits for paying out incentives. The credits will change over time based on the market share its FL model occupies. A player can choose not to join any Federation and just train models with their local dataset, or participate in a Federation. The process for joining a Federation involves three different stages: 1) Bidding, 2) FL model training and 3) Profit-sharing. In the bidding stage, participants choose bid to join a given Federation with his stated resource quantity, data quality, data quantity and payment. In the FL model training stage, the game simulates the training of the FL model based on the

participants’ bids. In the Profit-sharing stage, the Federation delivers payoffs to each participant following the incentive scheme it adopted before transitioning to the next bidding stage. *FedGame* supports the following incentive schemes [Yu *et al.*, 2020]:

- Linear: a participant’s share of the total payoff is proportional to the usefulness of its contributed data;
- Equal: the federation profit is equally divided among its participants;
- Individual: a participant’s share of the total payoff is proportional to its marginal contribution to the federation profit;
- Union: participant *i*’s share of the total payoff follows the Labour Union game payoff scheme and is proportional to the marginal effect on the FL model if *i* were to be removed;
- Shapley: the federation revenue is shared among participants according to their Shapley values.

At the same time, the system variable values which make up the context within which the players make decisions are also recorded to support further analysis of participant behaviours.

### 4 System Settings

The game system is configured using a text file that follows the XML format. Specified game settings, such as the number of players, types of Federations can be adjusted in *FedGame* through this configuration file. This facilitate game designers to vary the FL environment the players are exposed to. Besides the environmental variables, designers can adjust the time for FL model training, and the time taken for each round. Modification of these variables will allow for a shorter or longer game duration to influence participants’ behaviour.

### 5 Visualization

Figure 2 shows a screenshot for a player of the *FedGame* system. The game visualizes information including Federation information, game session overview, human player’s statistics, and game round summary to facilitate decision-making. It provides a continuous real-time view of the participants’ data quality, quantity, change of market share, profit/loss, and Federations’ participants. This simulates the information available to a sophisticated business joining FL to help researchers study possible reactions to given incentive mechanisms. A video and of the game system are available online<sup>1</sup>.

### 6 Conclusions and Future Work

The proposed *FedGame* system is, to the best of our knowledge, the first game for studying participants’ reactions under various incentive mechanisms in federated learning scenarios. Data collected can be used to analyse behaviour patterns exhibited by human players, and inform future FL incentive mechanism design research. In the future, we plan to extend the game with more complex processes and parameters to simulate Vertical Federated Learning (VFL) and Federated Transfer Learning (FTL) situations [Yang *et al.*, 2019b].




Figure 2: An example *FedGame* interface.

<sup>1</sup><https://youtu.be/UhAMVx8SOE8>

## Acknowledgments

This research is supported, in part, by Nanyang Assistant Professorship (NAP); NTU-WeBank JRI (NWJ-2019-007); AISG-GC-2019-003; NRF-NRFI05-2019-0002; and NTU-SDU-CFAIR (NSC-2019-011).

## References

- [Gollapudi *et al.*, 2017] Sreenivas Gollapudi, Kostas Kollias, Debmalaya Panigrahi, and Venetia Pliatsika. Profit sharing and efficiency in utility games. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA'17)*, 2017.
- [Jia *et al.*, 2019] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gurel, Bo Li, Ce Zhang, Dawn Song, and Costas Spanos. Towards efficient data valuation based on the shapley value. *CoRR, arXiv:1902.10275*, 2019.
- [Kairouz *et al.*, 2019] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G.L. D'Oliveira, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adriá Gascón, Badi Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrede Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. In *CoRR*, page arXiv:1912.04977, 2019.
- [McMahan *et al.*, 2016] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, et al. Communication-efficient learning of deep networks from decentralized data. *CoRR, arXiv:1602.05629*, 2016.
- [Pan *et al.*, 2009] Li Pan, Xiangxu Meng, Zhiqi Shen, and Han Yu. A reputation pattern for service oriented computing. In *Proceedings of the 7th International Conference on Information, Communications and Signal Processing (ICICS'09)*, pages 1–5, 2009.
- [Pan *et al.*, 2016] Zhengxiang Pan, Han Yu, Chunyan Miao, and Cyril Leung. Efficient collaborative crowdsourcing. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*, pages 4248–4249, 2016.
- [Shen *et al.*, 2011] Zhiqi Shen, Han Yu, Chunyan Miao, and Jianshu Weng. Trust-based web service selection in virtual communities. *Web Intelligence and Agent Systems*, 9(3):227–238, 2011.
- [Wu *et al.*, 2013] Qiong Wu, Xiaogang Han, Han Yu, Zhiqi Shen, and Chunyan Miao. The innovative application of learning companions in Virtual Singapura. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'13)*, pages 1171–1172, 2013.
- [Yang *et al.*, 2017] Shuo Yang, Fan Wu, Shaojie Tang, Xiaofeng Gao, Bo Yang, and Guihai Chen. On designing data quality-aware truth estimation and surplus sharing method for mobile crowdsensing. *IEEE Journal on Selected Areas in Communications*, 35(4):832–847, 2017.
- [Yang *et al.*, 2019a] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- [Yang *et al.*, 2019b] Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, and Han Yu. *Federated Learning*. Morgan & Claypool Publishers, 2019.
- [Yu *et al.*, 2010] Han Yu, Yundong Cai, Zhiqi Shen, Xuehong Tao, and Chunyan Miao. Agents as intelligent user interfaces for the net generation. In *Proceedings of the 15th International Conference on Intelligent User Interfaces (IUI'10)*, pages 429–430, 2010.
- [Yu *et al.*, 2011] Han Yu, Zhiqi Shen, Chunyan Miao, and Ah-Hwee Tan. A simple curious agent to help people be curious. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'11)*, pages 1159–1160, 2011.
- [Yu *et al.*, 2017] Han Yu, Zhiqi Shen, Chunyan Miao, Cyril Leung, Yiqiang Chen, Simon Fauvel, Jun Lin, Lizhen Cui, Zhengxiang Pan, and Qiang Yang. A dataset of human decision-making in teamwork management. *Scientific Data*, 4(160127):doi:10.1038/sdata.2016.127, 2017.
- [Yu *et al.*, 2018] Han Yu, Zhiqi Shen, Chunyan Miao, Cyril Leung, Victor R. Lesser, and Qiang Yang. Building ethics into artificial intelligence. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*, pages 5527–5533, 2018.
- [Yu *et al.*, 2020] Han Yu, Zelei Liu, Yang Liu, Tianjian Chen, Mingshu Cong, Xi Weng, Dusit Niyato, and Qiang Yang. A fairness-aware incentive scheme for federated learning. In *Proceedings of the 3rd AAAI/ACM Conference on AI, Ethics, and Society (AIES-20)*, pages 393–399, 2020.