

Resource-Aware Cooperation in Federated Learning

Manuel Röder^{1,2}, Fabian Geiger³ and Frank-Michael Schleif¹ *

1 - Technical UAS Würzburg-Schweinfurt, Fac. of Computer Science, Würzburg, DE

2 - Bielefeld University, Faculty of Technology, Bielefeld, DE

3 - JMU Würzburg, Fac. of Business Management and Economics, Würzburg, DE

Abstract. We present a novel Federated Learning framework, **FedT4T**, that systematically evaluates utility-driven client strategies under resource constraints. Recognizing the significant challenges in practical distributed learning environments, such as limited resources and non-cooperative behaviors, we model client interactions using the Iterated Prisoner’s Dilemma. Our framework enables clients to adapt their decision rules based on prior interactions and available resources, optimizing both individual utility and collective contribution to solve a global learning task. We apply **FedT4T** to a Federated Learning benchmark classification task and explore the dynamics of cooperation between clients driven by common strategies from cooperation theory under the impact of varying resource availability. The code is publicly available at <https://github.com/cairo-thws/FedT4T>.

1 Introduction

In the idealized vision of *Federated Learning* (**FL**) introduced by McMahan [1], each client within a decentralized network functions as a collaborative participant, utilizing its private data to train local models and subsequently transmit updated model parameters to contribute to a shared global model. This collaborative approach promises to leverage the potential of distributed data in a way that respects privacy and facilitates shared learning across a range of devices. This ideal presumes both uniform willingness and resource availability among clients, a condition seldom met in practical FL scenarios. In reality, resource limitations, communication overhead, and strategic client behaviors undermine the foundational assumptions of cooperative, resource-rich participation [2].

To address the complexities of client cooperation and resource allocation in FL environments, we draw upon the conceptual framework of the *Iterated Prisoner’s Dilemma* (**IPD**). The IPD, a well-established model in cooperation theory [3], provides a foundation for analyzing strategic interactions in which participants choose between cooperative and non-cooperative behaviors over repeated encounters. In the context of FL, the IPD framework allows us to interpret client interactions as a series of strategic decisions regarding resource allocation and model contribution. Clients in an FL setting must decide whether to cooperate — expending scarce resources to contribute to the global model — or to defect, potentially conserving resources but providing less support for the distributed learning task. In this work, we systematically evaluate various IPD strategies within FL environments in a simulated round-robin tournament to understand

*MR is supported by the Würzburg Center for Artificial Intelligence and Robotics (CAIRO), THWS ProPere and the European Regional Development Fund (ERDF), FKZ:2404-003-1.2

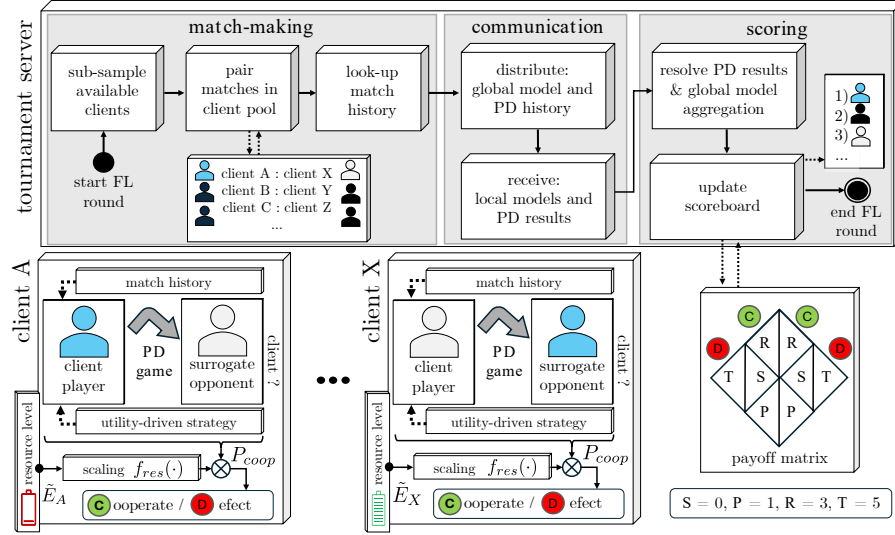


Fig. 1: Overview of FedT4T: A round of FL managed by a central tournament server.

their impact on distributed model performance as outlined in Fig. 1. Specifically, we investigate how different strategic choices by FL clients influence their ability to contribute to the shared learning objective under conditions of varying resource availability. Using a payoff scoring scheme, we quantify the outcomes of these interactions, assigning scores to clients based on their strategies and levels of cooperation. This approach enables the identification of the trade-offs between cooperation and resource conservation across a range of scenarios involving client-specific resource utilization. By comparing FL client behaviors across varying levels of resource availability, we aim to identify strategies that optimize both individual client performance and collective model quality. Our analysis highlights the effectiveness of certain IPD strategies for balancing resource constraints with the goal of sustaining effective FL. Through this paper, we provide a novel perspective on client strategy optimization in resource-limited FL environments, offering insights into the role of cooperative behaviors in distributed learning systems. *Recent works* have addressed game-theoretical approaches within distributed learning settings [4, 5]. However, these approaches establish central instances that compete against clients in IPD games directly and they do not consider the important impact of client-side resource availability on the participation policy.

The **main contributions** of this work are threefold. First, we conceptualize the FL environment as an infinite IPD (see Fig. 1, top), where the FL tournament server is responsible for recording client decisions and facilitating encounters between client pairs in each FL round. Second, we propose a novel approach where clients locally perform a single turn of the PD against a projected surrogate opponent (see Fig. 1, bottom), adapting their utility-driven strategy based on prior interaction. Finally, clients determine their participation decisions by incorporating their available resources—a unique feature that enables an assessment of FL training dynamics under more realistic, resource-constrained conditions.

2 Methodology

2.1 Setup and Prerequisites

We conceptualize a round of FL as a series of randomly sampled interactions between two FL clients, structured within a round-robin tournament. Each client has the option to either cooperate (C) by training on its locally available data and submitting updated model weights for aggregation, or defect (D) by conserving resources and skipping the local training process. As illustrated in Fig. 1, the tournament server is responsible for matchmaking and, when available, providing each client with the outcome of its most recent encounter against a selected surrogate opponent (*memory*). Additionally, the server manages the collection and aggregation of model weights, as well as the resolution and bookkeeping of scores derived from the PD results. Each client’s payoff is determined by both its own decision and that of its surrogate opponent, following the 2×2 payoff matrix in Fig. 1 (bottom right). Here, R (Reward) represents the payoff when both clients cooperate, T (Temptation) benefits a defecting client when the opponent cooperates, S (Sucker’s payoff) applies to a cooperating client when the opponent defects, and P (Punishment) is assigned when both clients defect. The payoffs adhere to the conditions $T > R > P > S$ and $2R > T + S$, ensuring that mutual cooperation leads to a more favorable collective outcome compared to mutual defection or unilateral cooperation. As detailed in Sec. 2.2 and without loss of generality, this paper focuses exclusively on FL clients that adopt *memory-one* (M1) strategies [3]. Accordingly, let p_1 , p_2 , p_3 , and p_4 denote the probabilities of a client choosing to cooperate, conditioned on the joint action pairs from the previous round: (CC), (CD), (DC), and (DD), respectively. These probabilities reflect a utility-driven decision-making process – the stochastic decision rule vector for client i in FL round t is thus given by $\mathbf{p}_i^t = [p_1, p_2, p_3, p_4] \in \mathbb{R}_{[0,1]}^4$ where $[1, 1, 1, 1]$ represents the standard FL strategy of unconditional cooperation [6].

2.2 Client Memory Depth and Cooperation through Repetition

When modeling FL interactions, we opt to represent the game as an *infinitely repeated PD* rather than a finite version. In finite repeated games, backward induction implies that defection becomes the dominant strategy in the final rounds, as both players anticipate the end of their interactions. This discourages cooperation, as FL clients would have an incentive to defect once they recognize the game has a predefined endpoint. In contrast, an infinitely repeated setting eliminates a fixed terminal payoff, allowing client payoffs to be defined based on average or discounted rewards. This framework aligns naturally with the FL environment, where the total number of rounds remains undisclosed to clients and is influenced by random client sub-sampling in each training round to prevent pattern exploitation. As a result, clients cannot predict a finite number of interactions, maintaining the conditions necessary for strategies that promote sustained cooperation and reciprocity.

For efficient decision-making, our framework focuses specifically on M1 strategies, where each client’s next action depends solely on the outcome of the most recent round: Press and Dyson’s findings in [7] demonstrate that, in repeated

games, long-memory strategies provide no advantage over short-memory ones, as players can achieve favorable outcomes by considering only the previous interaction. This approach not only minimizes computational complexity but also aligns with the resource constraints of FL environments, where lightweight decision-making is essential. Therefore, we exclude long-memory agents, such as those utilizing LSTM-based reinforcement learning, as well as adaptive decision rules that evolve over time. Instead, we prioritize simpler, more efficient strategies that are better suited to the immediate-response requirements of FL trainings.

2.3 Resource-Aware Decision-Making for Federated Learning Clients

In practical FL deployments, clients often have varying resource constraints, such as limited battery life, computational capacity, or network bandwidth. To address these variations, we incorporate resource awareness into the client participation strategy by adapting concepts from the IPD. An advantage of integrating resource-aware scaling into M1 strategies in FL is that it prevents clients from unilaterally controlling or manipulating the game mechanics through *zero-determinant* (**ZD**) strategies [7]. By individually adjusting cooperation probabilities based on resource availability, the precise linear relationships required for ZD strategies are disrupted, eliminating the possibility for any client to enforce unilateral payoff control.

To roll-out resource awareness to FL clients, we introduce a resource scaling function for adjusting the overall cooperation probability P_{coop}^t of client i in round t (omitted in the following for better readability), effectively blending resource considerations into the decision-making process: Let us first denote the current resource availability of client i by E_i , $E_i \in [0, E_{\text{Max}}]$, where E_{Max} represents the maximum possible resource level. We subsequently normalize E_i using $\tilde{E}_i = E_i/E_{\text{Max}}$ and define a resource scaling function $f_{\text{res}}(\tilde{E}_i)$ that consumes and adapts the normalized resource level. To ensure robust probability of participa-




Fig. 2: Synergy Threshold Function with hyper-parameter $\gamma = \{4, 8, 12\}$ and distinct zones for low, moderate, high and full resource levels.

tion in regimes with higher resource availability, we propose a *Synergy Threshold Function* for scaling, illustrated in details in Fig. 2, defined as

$$f_{\text{res}}(\tilde{E}_i) = \frac{1}{2} \left(1 + \tanh \left(\gamma (\tilde{E}_i - E_{\text{low}}) \right) \right), \quad (1)$$

where $\gamma > 0$ controls the steepness around the critical low-resource threshold E_{Low} , offering a smooth and adjustable transition in cooperation prob-

ability based on resource levels and encouraging clients to participate even when resources are low. To subsequently obtain the adjusted overall cooperation probabilities, each client updates its utility-driven cooperation probability P_{coop} using the result of the deployed resource scaling function defined as $P'_{\text{coop}} = f_{\text{res}}(\tilde{E}_i) \times P_{\text{coop}}$. Applied to the stochastic decision rules of our framework described in Sec. 2.1, the resource-aware adjustment can be expressed as

$$P'_{\text{coop}} = f_{\text{res}}(\tilde{E}_i) \times p_{\text{out}}, \quad (2)$$

where $p_{\text{CC}} = p_1$, $p_{\text{CD}} = p_2$, $p_{\text{DC}} = p_3$, $p_{\text{DD}} = p_4$

and *out* referring to the outcome of the last encounter. Each FL client will decide to cooperate or defect based on P'_{coop} , which now factors in both utility-driven choices obtained from M1 strategies and current resource availability. The modified decision rule vector governs whether client i participates in training or not by evaluating $U(0,1) < P'_{\text{coop}}$, where $U(0,1)$ represents a random value sampled from a uniform distribution.

3 Experiment and Discussion

In our experimental evaluation of cooperative behavior within the **FedT4T** framework, we observe a set of clients solving a benchmark FL classification task conducted under non-IID data distribution setting [8]. We implement the tournament server and the IPD clients as outlined in Fig. 1 using Flower [9], provide the utility strategies from the Axelrod library [10], integrate resource-awareness as described in Sec. 2 and employ the Synergy Threshold function ($\gamma = 8$) as proposed in Eq. (1). We provide cooperation statistics of eight distinct resource-aware FL client strategies, driven by basic M1 decision rules over 250 FL rounds and report the results in Table 1. We supportively present plots of the cu-

Table 1: **FedT4T** evaluation results of resource-aware FL clients, driven by distinct M1 decision rules with *nice* properties [3], sorted by average payoff.

IPD Strategy Name	Memory One Parameter	Total Payoff	Average Payoff	Defections	Cooperations	Cooperation Rate (%)
Grim (GRIM)	[1, 0, 0, 0]	658	2.64	101	148	59.44
Tit for Tat (TFT)	[1, 0, 1, 0]	622	2.50	94	155	62.25
FedT4T Custom (Contributor)	[.9, .5, .5, .9]	615	2.47	121	128	48.59
Generous TFT (GTFT)	[1, .33, 1, .33]*	605	2.43	84	165	66.26
Firm But Fair (FirmButFair)	[1, 0, 1, .66]	605	2.43	71	178	71.49
Win Stay-Lose Shift (WLS)	[1, 0, 0, 1]	602	2.42	79	170	68.27
Forgiving TFT (FTFT)	[1, .75, 1, .75]	593	2.38	54	195	78.31
Cooperator (Cu)	[1, 1, 1, 1]	571	2.29	44	205	82.33

* $[1, \min(1 - \frac{r-p}{r-s}, \frac{r-p}{r-s}), 1, \min(1 - \frac{r-p}{r-s}, \frac{r-p}{r-s})]$

mulative client cooperations, tracked over all server rounds, in Fig. 3. Resource availability declines stepwise below E_{High} in round 50, E_{Moderate} in round 100 and E_{Low} in round 150, respectively, simulating an environment with progressively constrained computational and communication capacities: In the full-resource phase, FL clients exhibit high levels of cooperation, reflecting Pareto-efficient outcomes where all clients benefit mutually. However, as resources degrade, cooperative dynamics shift towards Nash equilibria, where self-interest drives decision-making and defections increase [3]. The experiment unveils another interesting finding: while the standard FL strategy **Cu** achieves high cooperation




Fig. 3: Cumulative cooperation scores of FedT4T clients under stepwise resource reduction, with focus area on the threshold to the low-resource sector. Best viewed in color. rates (Fig.3), its inability to adapt to resource constraints results in poor payoff maximization (Tab. 1). This highlights its unsuitability for resource-aware environments, where dynamic and conditional strategies prove to be more effective.

4 Conclusion

In this work, we introduced a novel perspective on FL by formulating the distributed training process as an IPD tournament, where client decisions are driven by their utility policy and local resource availability. First experimental results showed that FedT4T is a powerful tool for the application of cooperation theory in the analysis of FL trainings. *Future work* includes the application of incentive mechanisms on FL clients based on IPD metrics, the study of the impact of their cooperation behavior on the convergence of the global FL model, and the evaluation of a broader range of resource scaling functions.

References

- [1] B. McMahan et al. Communication-efficient learning of deep networks from decentralized data. In *Proc. 20th AISTAT*, volume 54, pages 1273–1282, April 2017.
- [2] M. Röder et al. Crossing Domain Borders with Federated Few-Shot Adaptation. In *Proc. 13th ICPRAM*, pages 511–521, 2024.
- [3] R. Axelrod et al. *The evolution of cooperation*. Basic Books, New York, NY, 2006.
- [4] A. Akbay et al. Distributed Learning with Strategic Users: A Repeated Game Approach. *AAAI*, 36(6):5976–5983, June 2022.
- [5] X. Zhang et al. A game-theoretic framework for privacy-preserving federated learning. *ACM Trans. Intell. Syst. Technol.*, 15(3), May 2024.
- [6] S. Kuhn. Prisoner’s Dilemma. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2024 edition, 2024.
- [7] W. Press et al. Iterated Prisoner’s Dilemma contains strategies that dominate any evolutionary opponent. *Proc. Natl. Acad. Sci. U.S.A.*, 109(26):10409–10413, June 2012.
- [8] Q. Li et al. Federated learning on non-iid data silos: An experimental study. *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 965–978, 2021.
- [9] D. Beutel et al. Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390*, 2020.
- [10] V. Knight et al. An open reproducible framework for the study of the iterated prisoner’s dilemma. *Journal of Open Research Software*, 4, 04 2016.