Federated Learning by Utility-Constrained Stochastic Aggregation for Improving Rational Participation (NeurIPS, 2026)

M Yashwanth, Arunabh Singh, Ashok Nayak, Saikiran Bulusu, Anirban Chakraborty

Sunday 15th November, 2026

Motivation

- In Federated Learning, its commonly assumed that the clients are passive in nature, they participate in federation when sampled by server.
- This passive participation is reasonable in cross-device FL, where the the edge-device itself may not hold sufficient data. This assumption can potentially fail in the cross-silo scenario.
- In the cross-silo setup, which is typically a institutional setup, the clients can be rational they will not participate in federation unless it benefits them.

Monetary Vs Non-Monetary Incentives

- Server typically incentivize clients to participate in Federation, this incentive very commonly involves monetary aspect.

- In this work, we question what if such a payment channel doesn't exist, or the clients themselves are not interested in monetary incentives but rather interested in the quality of the global model.

- Especially they will contribute to federation, if they satisfy Individual Rationality constraint

$$acc(\boldsymbol{\theta}^t, \mathcal{D}_k^{\text{val}}) > thr_k,$$

where $\boldsymbol{\theta}^t = \sum_{k=1}^{N} \alpha_k^{t-1} \boldsymbol{\theta}_k^{t-1}$. and $thr_k = acc(\boldsymbol{\theta}_k^{t-1}, \mathcal{D}_k^{\text{val}})$, which implies global model is better than local model.

- The threshold can also be set to a convex combination of local and global models $thr_k = \alpha * acc(\boldsymbol{\theta}_k^{t-1}, \mathcal{D}_k^{\text{val}}) + (1 - \alpha) * acc(\boldsymbol{\theta}^{t-1}, \mathcal{D}_k^{\text{val}})$.

Monetary Vs Non-Monetary Incentives

- We cannot assume clients to participate with such strict rationality check, as there is no control with the server to provide a model which is better for the client.

- We now develop the non-monetary machinery where the clients can do a non-monetary transactions with the server or reason with the server on its incentives.

- We propose that server shares the client's aggregation weight with each client used in the global model.

Utility Function Design

- The main idea is to design the utilty function of the client $u_k(.)$ as a function of aggregation weight r_k.

- The client will now participate if the aggregation weight provided by the server satisfies the following

$$u_k(r_k; \zeta_k) > \gamma_k,$$

  where ζ_k is the parameters that define the utility.

- The utility here can be considered to be accuracy or the negative loss on the held-out client data.

Utility Function Design

- Since the server shares the aggregation weight with the client, the client can form the aggregated model without its own contribution as

$$\boldsymbol{\theta}_{-k}^{t} = \frac{\boldsymbol{\theta}^{t} - \alpha_k^{t-1}\boldsymbol{\theta}_k^{t-1}}{1 - \alpha_k^{t-1}}$$

- Once we get the held-out global model, the client will linearly combine its locally trained model $\tilde{\boldsymbol{\theta}}_k(r)$ with an aggregation parameter r as $\tilde{\boldsymbol{\theta}}_k(r) = (1-r)\boldsymbol{\theta}_{-k}^{t} + r\boldsymbol{\theta}_k^{t}$, Now $\tilde{\boldsymbol{\theta}}_k(r)$ is interpreted as a global model with client k aggregation weight as r.

- We evaluate the accuracy at this interpolated global model as

$$\hat{u}_k(r) = \mathop{\mathbb{E}}_{(\mathbf{x},y)\sim\mathcal{D}_k^{val}}[acc(h(\mathbf{x};\tilde{\boldsymbol{\theta}}_k(r)), y)]$$

- We finally fit this data to get the parametrized concave utility that can be communicated with the server.

$$\zeta_k := \arg\min_{\tilde{\zeta}_k} \mathbb{E}_{r\sim\mathcal{U}(0,1)}\left[\|u_k(r;\tilde{\zeta}_k) - \hat{u}_k(r)\|_2^2\right].$$

On the concave modelling of the utility

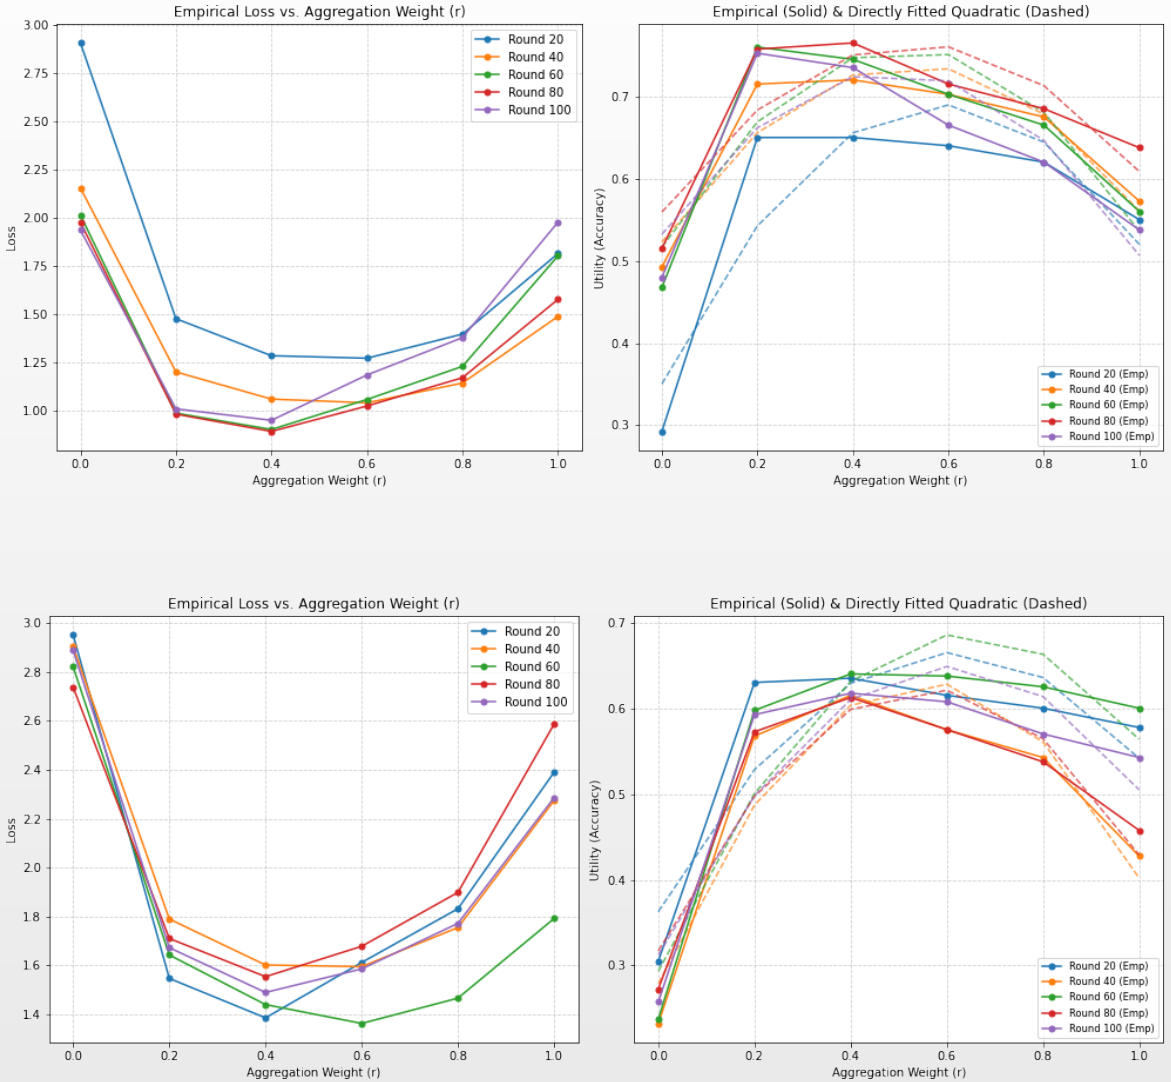

Figure 1: Representative empirical loss and utility profiles across two different clients, rounds, on the CIFAR-100 dataset. Across settings, loss curves exhibit convex-like behavior along the interpolation path, while utility curves remain smooth and

Designing the Server Optimization Problem

- Each client will request aggregation weight such that, $u_k(r_k; \zeta_k) > \gamma_k$
- Rather than optimizing over a single, static weight vector r_k, the server generates a set of n_s candidate strategies $\{r_k^j\}_{k=1}^N$ for $j = 1, \ldots, n_s$.
- The influence of these strategies is determined by a simplex vector $\mathbf{s} = \{s_1, s_2, \ldots, s_{n_s}\}$, where s_j represents the probability of selecting the jth strategy.
- This vector $\mathbf{s}$ is sampled from the Dirichlet distribution with concentration δ.
- The server communicates to client k the expected aggregation weight $\alpha_k = \sum_i s_i r_k^i$.

Designing the Server Optimization Problem

$$
\begin{aligned}
\underset{\mathbf{t},\mathbf{R}}{\text{minimize}} \quad & \frac{1}{2}\|\mathbf{t}\|_2^2 \\
\text{subject to} \quad & \mathbf{R}\mathbf{1}_N = \mathbf{1}_{n_s}, \mathbf{1}_{n_s}^T\mathbf{R} = \frac{n_s}{N}\mathbf{1}_N^T, \mathbf{r}_k \succeq 0, \mathbf{t} \succeq 0 \\
& \sum_{i=1}^{n_s} s_i u_k(r_k^i) \geq \epsilon + \gamma_k - t_k, \forall k \in [N]
\end{aligned} \tag{1}
$$

- Here $\mathbf{R}$ is the matrix of size $n_s \times N$, the aggregation weight matrix. t_k is the slack variable and $\epsilon > 0$ is the saftey margin.
- Key thing to note is we optimize for lower bound. Since the utilities are concave, we have the following by Jensen's inequality, $u_k(\sum_{i=1}^{n_s} s_i r_k^i) \geq \sum_{i=1}^{n_s} s_i u_k(r_k^i)$.

Solving the Optimization

Proposition

Suppose $\epsilon = 0$, *and all utilities* $u_k(.)$ *are concave then the solution to optimization 1 satisfies either* $\sum_{i=1}^{n_s} s_i u_k(r_k^i) = \gamma_k - t_k$, *with* $t_k > 0$ *or* $\sum_{i=1}^{n_s} s_i u_k(r_k^i) \geq \gamma_k$, *with* $t_k = 0$. *In the former case, the constraint implies* $u_k(r_k^i) = \frac{s_i}{\sum_{j=1}^{n_s} s_j^2}(\gamma_k - t_k)$. *If we further assume that* $u_k(.)$ *is strictly concave, the optimal solution* $\mathbf{R}^*, \mathbf{t}^*$ *is unique, if the former holds* $\forall k \in [N]$.

Proof Sketch.

The proof follows by analyzing the KKT conditions. The details are provided in the paper. The generality of the result holds independent of ϵ as it can always be absorbed into the γ_k. □

The Jensen Gap

- We analyze the induced Jensen's gap associated with client k,

$$\mathcal{J}_k(\mathbf{s}) = u_k(\sum_{i=1}^{n_s} s_i r_k^i) - \sum_{i=1}^{n_s} s_i u_k(r_k^i), \tag{2}$$

  as a function of s.

- This is a key quantity, beacuse if the Jensen gap can compensate for the decrease due to slack, the client will participate, and the key is to control the Jesnen Gap.

Controlling the Jensen Gap with Dirichlet Concentration δ

Proposition

Let the expected Jensen gap for client k be

$$\bar{\mathcal{J}}_k(\delta) := \mathbb{E}_{\mathbf{s}\sim\mathrm{Dir}(\delta)}[\mathcal{J}_k(\mathbf{s})], \qquad \delta > 0,$$

with instantaneous gap

$$\mathcal{J}_k(\mathbf{s}) := u_k\Big(\sum_{i=1}^{n_s} s_i r_k^i\Big) - \sum_{i=1}^{n_s} s_i\, u_k(r_k^i),$$

where u_k is strictly concave and the weights r_k^i solve Optimization 1. Then:

- *$\bar{\mathcal{J}}_k(\delta) \to 0$ as $\delta \to 0$ and as $\delta \to \infty$, and is small for extreme finite values $\delta \le \delta_{\min}$ or $\delta \ge \delta_{\max}$;*
- *$\bar{\mathcal{J}}_k(\delta)$ is continuous on $(\delta_{\min}, \delta_{\max})$ and attains its maximum there.*

Controlling the Jensen Gap: Proof Sketch

Proof Sketch.

1. **Measurability.** $\mathcal{J}_k(\mathbf{s})$ is a measurable function of $\mathbf{s}$, so the expectation $\bar{\mathcal{J}}_k(\delta)$ is well defined.
2. **Limiting regimes.** Analyzing $\delta \ll 1$ and $\delta \gg 1$: the Dirichlet mass concentrates on a vertex ($\delta \to 0$) or at the center ($\delta \to \infty$); in both the gap collapses, so $\bar{\mathcal{J}}_k(\delta) \to 0$.
3. **Continuity.** By the Dominated Convergence Theorem, $\bar{\mathcal{J}}_k(\delta)$ is continuous in δ on $(\delta_{\min}, \delta_{\max})$.
4. **Interior maximum.** The gap is strictly positive on the interior; continuity on the interval then forces the maximum to be attained at some $\delta^\star \in (\delta_{\min}, \delta_{\max})$.

Full proof in the paper. □

Empirical Verification of Proposition

Table 1: Empirical verification of Proposition. We report the per-round total Jensen's gap $\sum_k \mathcal{J}_k^t$ measured over 100 communication rounds on CIFAR-10, summarized three ways: the mean over all rounds, the mean over the 50 rounds with the largest gap (Top-50), and the mean over the 25 rounds with the largest gap (Top-25). The Top-50 and Top-25 columns characterize the rounds in which the surplus mechanism is most active, typically rounds in which a substantial fraction of clients have binding IR constraints. The non-monotonic dependence on δ predicted by Proposition 12 is observed across all three columns: the gap vanishes at extreme concentrations ($\delta = 0.01$ and $\delta = 100$) and peaks at moderate values $\delta \in [0.5, 1.0]$.

Dirichlet Concentration (δ)	**Mean**	**Top-50 Mean**	**Top-25 Mean**
0.01	0.20	0.40	0.81
0.5	2.48	4.10	5.32
1.0	2.77	4.26	5.18
10	1.19	1.60	1.82
100	0.31	0.48	0.67

Empirical Verification of Multiple Strategies

Table 2: Server Accuracy between single and multiple strategies

Strategies	**CIFAR-10** (ACC (%)/ Participation %)	**CIFAR-100** (ACC % / Participation %)
$n_s = 1$	63.80 (29)	36.80 (37)
$n_s = 5$	66.89 (53)	42.50 (53)

Convergence Decomposition (Assumptions)

- The standard non-convex FL assumptions used in partial-participation analyses: each f_k is L-smooth.
- Stochastic gradients are unbiased with variance at most σ^2, client heterogeneity satisfies $\|\nabla f_k(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta})\|^2 \leq \sigma_g^2$.
- local gradients are bounded as $\|\nabla f_k(\boldsymbol{\theta})\|^2 \leq G$
- stale delays satisfy $\nu_k^t \leq \nu_{\max}$.
- We assume that the participation set S_t is uniformly lower bounded, i.e., $1 \leq S \leq |S_t| \leq N$, where S is the minimum number of participating clients per round.

FedUCA Convergence: Error Decomposition

Proposition (Convergence Decomposition, informal)

Under smoothness, bounded variance, bounded gradients, and bounded stale-delay assumptions, FedUCA satisfies

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\|\nabla f(\boldsymbol{\theta}^{(t,0)})\|^2 \leq \mathcal{E}_{\text{init}} + \mathcal{E}_{\text{agg}} + \mathcal{E}_{\text{var}} + \mathcal{E}_{\text{drift}} + \mathcal{E}_{\text{sys}}.$$

- $\mathcal{E}_{\text{init}} = \mathcal{O}(1/\sqrt{T})$: optimization term.
- $\mathcal{E}_{\text{agg}} \propto \frac{1}{T}\sum_t \mathbb{E}\|\boldsymbol{\beta} - \boldsymbol{\alpha}^t\|^2$: aggregation discrepancy.
- $\mathcal{E}_{\text{var}}, \mathcal{E}_{\text{drift}}$: Vanishing stochastic variance, local and stale-update drift.
- $\mathcal{E}_{\text{sys}}$: system term with a non-vanishing bias $\mathcal{O}((N-S)\sigma^2)$.

Full statement and proof in the paper.

FedUCA Convergence: Proof Sketch

Proof Sketch.

1. **Condition on history H_t:** reported utilities, thresholds, participation, and aggregation weights become fixed for round t; the server-side optimization is abstracted into the per-round weights $\boldsymbol{\alpha}^t$ and cohort S_t.
2. **Aggregation penalty:** smoothness yields the discrepancy term $\mathcal{E}_{\text{agg}} \leq \mathbb{E}\|\boldsymbol{\beta} - \boldsymbol{\alpha}^t\|^2$. With $\boldsymbol{\beta} = \frac{1}{N}\mathbf{1}_N$,

$$\mathbb{E}\|\boldsymbol{\beta} - \boldsymbol{\alpha}^t\|^2 \leq \frac{n_s - 1}{N(n_s\delta + 1)} \;\Rightarrow\; \mathcal{E}_{\text{agg}} \leq \frac{C(n_s - 1)}{S(n_s\delta + 1)}.$$

3. **Drift terms:** active-client drift from standard local-SGD bounds; stale-client drift controlled by the delay bound $\nu_k^t \leq \nu_{\max}$ and the minimum cohort size S.
4. **Telescope** the descent inequality over T rounds to obtain the decomposition.

□

Experimental results

Table 3: Server Model performance on CIFAR-10, F-MNIST, and CIFAR-100 under different data splits. The data partition Dirichlet (0.5, 1.0) implies that training data is distributed according to the Dirichlet distribution based on the concentration parameter 0.5, and the validation data uses 1.0.

	CIFAR-10		F-MNIST		CIFAR-100	
	Dirichlet 0.5,1.0	Dirichlet 1.0,1.0	Dirichlet 0.5,1.0	Dirichlet 1.0,1.0	Dirichlet 0.5,1.0	Dirichlet 1.0,1.0
FedAvg	$16.63_{\pm 1.16}$	$18.31_{\pm 1.66}$	$53.26_{\pm 2.78}$	$54.64_{\pm 2.05}$	$04.05_{\pm 0.06}$	$03.63_{\pm 0.50}$
IncFL	$32.35_{\pm 0.96}$	$34.60_{\pm 0.16}$	$37.07_{\pm 0.20}$	$36.69_{\pm 0.04}$	$04.18_{\pm 0.41}$	$12.89_{\pm 0.94}$
FedProx	$17.44_{\pm 0.86}$	$17.63_{\pm 2.14}$	$54.68_{\pm 2.01}$	$53.79_{\pm 0.55}$	$06.35_{\pm 2.91}$	$04.06_{\pm 0.22}$
AAggFF-S	$16.13_{\pm 0.50}$	$17.90_{\pm 0.33}$	$52.84_{\pm 3.18}$	$68.92_{\pm 6.22}$	$5.60_{\pm 0.33}$	$04.53_{\pm 0.86}$
FedAvg-WS	$47.70_{\pm 0.47}$	$51.63_{\pm 0.71}$	$77.60_{\pm 0.57}$	$75.75_{\pm 0.23}$	$33.50_{\pm 0.26}$	$48.93_{\pm 0.05}$
FedUCA (Ours)	$\mathbf{66.89_{\pm 01.56}}$	$\mathbf{80.00_{\pm 0.69}}$	$\mathbf{83.18_{\pm 0.40}}$	$\mathbf{83.48_{\pm 0.85}}$	$\mathbf{42.50_{\pm 0.29}}$	$\mathbf{50.15_{\pm 0.32}}$

Participation

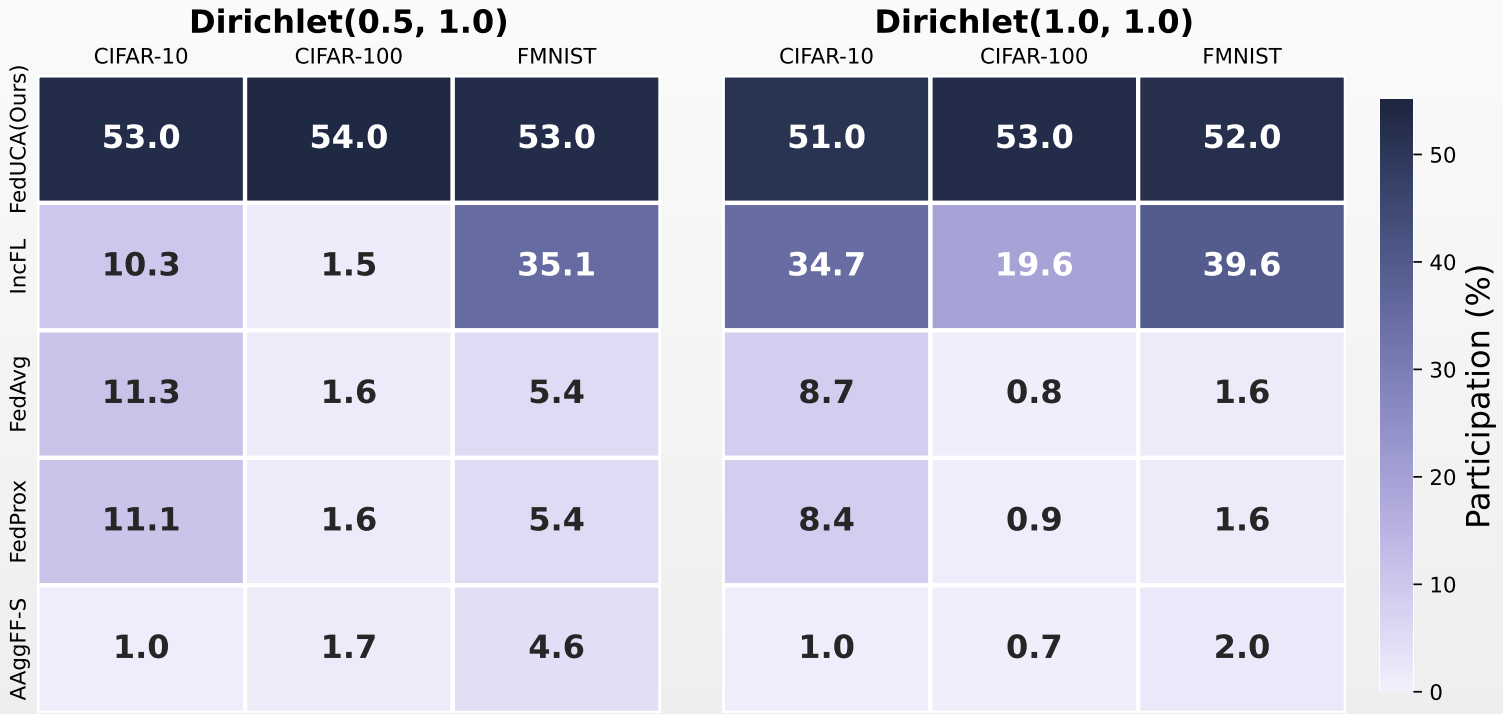

Figure 2: Comparison of participation rates across different FL methods on CIFAR-10, CIFAR-100, and FMNIST under varying degrees of data heterogeneity Dirichlet(0.5, 1.0) and Dirichlet(1.0, 1.0).

Summary

- We proposed a non-monetary framework for rational Clients in FL.
- We show that randomized aggregation is better than the deterministic one.
- We provide Convergence characterization for the endogenous participation and aggregation.
- Our empirical analysis show the model do attain convergence, it remains open to show what fixed-point the model converges to.
- How the clients can exploit utility functions, providing the Incentive compatibility guarantees are left to future work.