Redefining Contributions: Shapley-Driven Federated Learning

MOHAMED BIN ZAYED UNIVERSITY OF ARTIFICIAL INTELLIGENCE

IJCAI JEJU 2024

Nurbek Tastan, Samar Fares, Toluwani Aremu,
Samuel Horvath, Karthik Nandakumar

Problem Definition and Contribution

Goal: Assessment of class-specific contributions of participants in a federated learning setting, which aids in measuring the statistical heterogeneity.

Key Contributions:

- We introduce ShapFed, a novel method for precisely quantifying each participant's impact on the global model, including overall and class-specific contributions.
- Building upon our contribution assessment approach, we propose a new weighted aggregation method (ShapFed-WA) that outperforms the conventional federated averaging algorithm.
- To enhance collaborative fairness, we personalize server-to-client updates based on contributions, ensuring that substantial contributors receive better updates than those with minimal input.

Formulation

Problem: Standard cross-silo federated learning optimization problem

$$f^\star := \min_{w \in \mathbb{R}^d} \left[f(w) := \frac{1}{n} \sum_{i=1}^{n} f_i(w) \right], \qquad (1)$$
$$f_i(w) := \mathbb{E}_{\xi \sim \mathcal{D}_i} \left[F_i(w, \xi) \right]$$

Method

Overview of our proposed ShapFed algorithm:

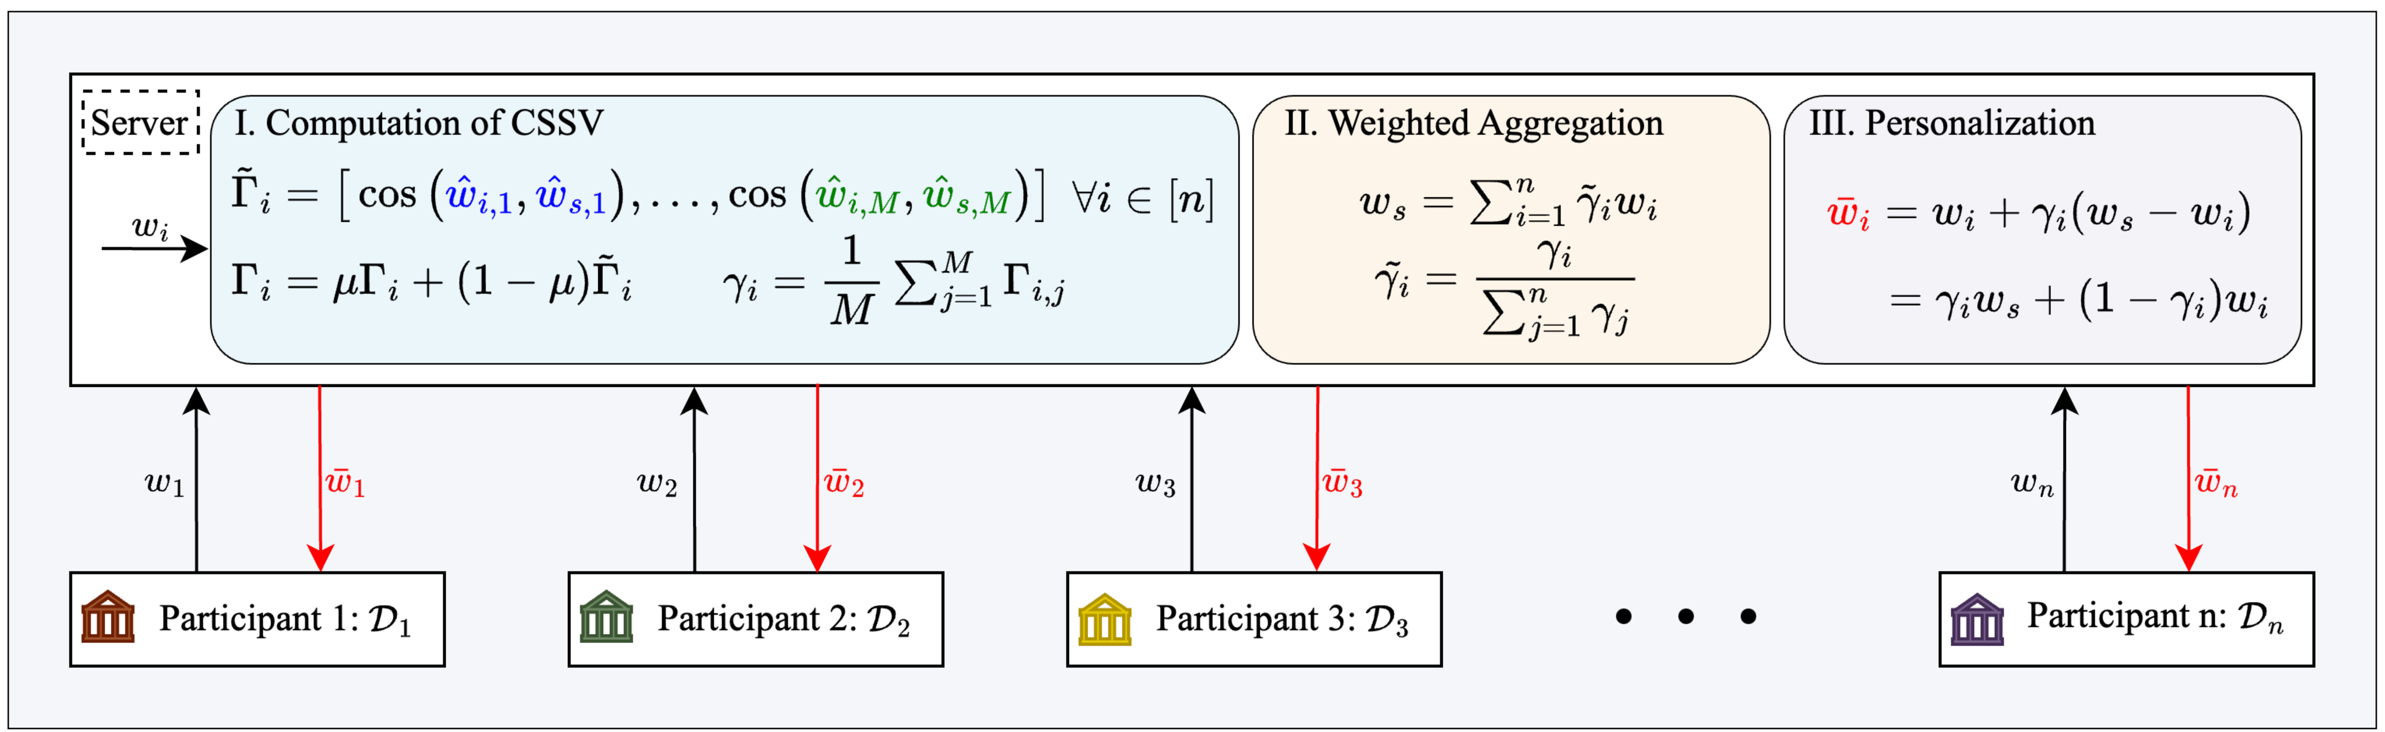

Each participant i transmits their locally computed iterates w_i to the server. The server then,

i. computes class-specific Shapley values (CSSVs) using the last layer parameters (gradients) $\hat{w}$,
ii. aggregates the weights by employing normalized contribution values $\tilde{\gamma}_i$ for each participant i,
iii. broadcasts the personalized weights $\bar{w}_i$ to each participant, using their individual, not-normalized contribution values γ_i.

Segment of the network utilized for evaluating class-wise contributions:

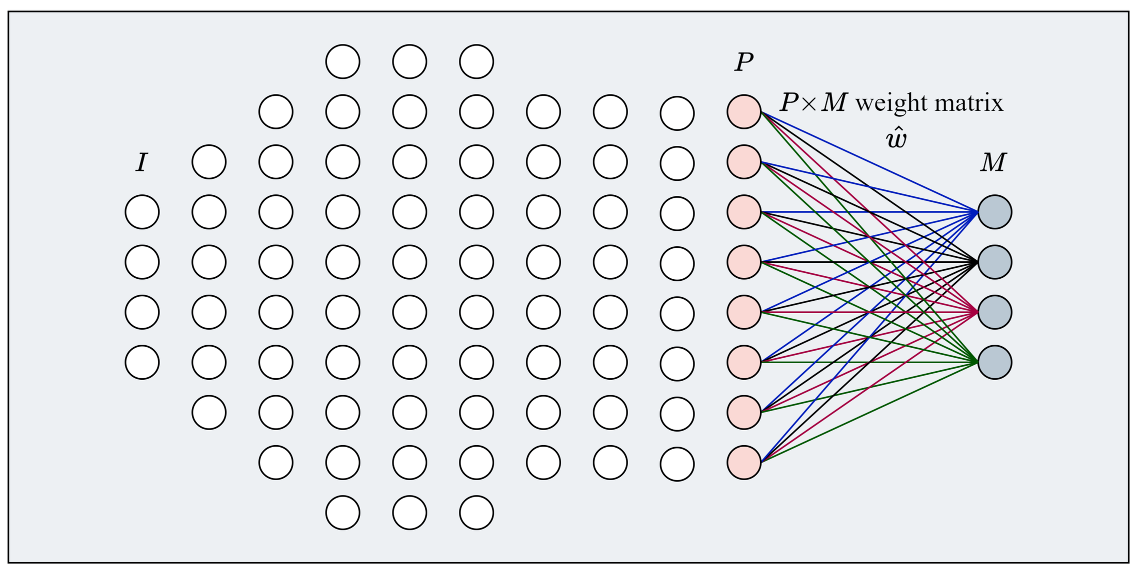

Weighted aggregation and personalization:

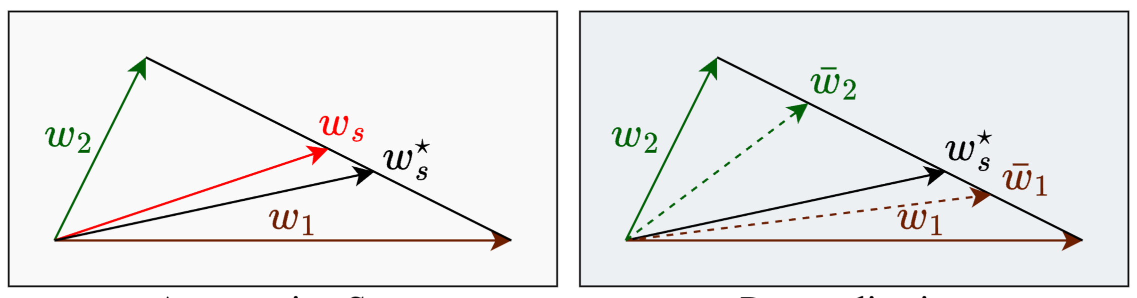

Aggregation Step   Personalization

Experiments & Results

Experimental Setup:

Dataset	Architecture	Batch size	Comm. rounds	Description
CIFAR-10	ResNet-34	128	50	60000 (10 classes)
Chest X-Ray	3-Conv & 3-FC	128	50	112120 (2 classes)
Fed-ISIC2019	EfficientNet_B0	32	200	23247 (8 classes)

Data Partitioning:

- **Imbalanced partitioning:** we use a custom function that relies on parameters x and y, where x determines the proportion of data points received by each of the y chosen participants. The remaining participants then share the remaining data among themselves.
- **Heterogeneous partitioning:**
  - **CIFAR-10:** class 1 is exclusively owned by participant 1, and the remaining 9 classes are partitioned equally among all participants.
  - **Chest X-Ray:** with 5 participants scenario, class 1: [40%, 30%, 20%, 10%, 0%] and class 2: [0%, 10%, 20%, 30%, 40%].

Optimizer: SGD with learning rate 0.01.

Weighted Aggregation:

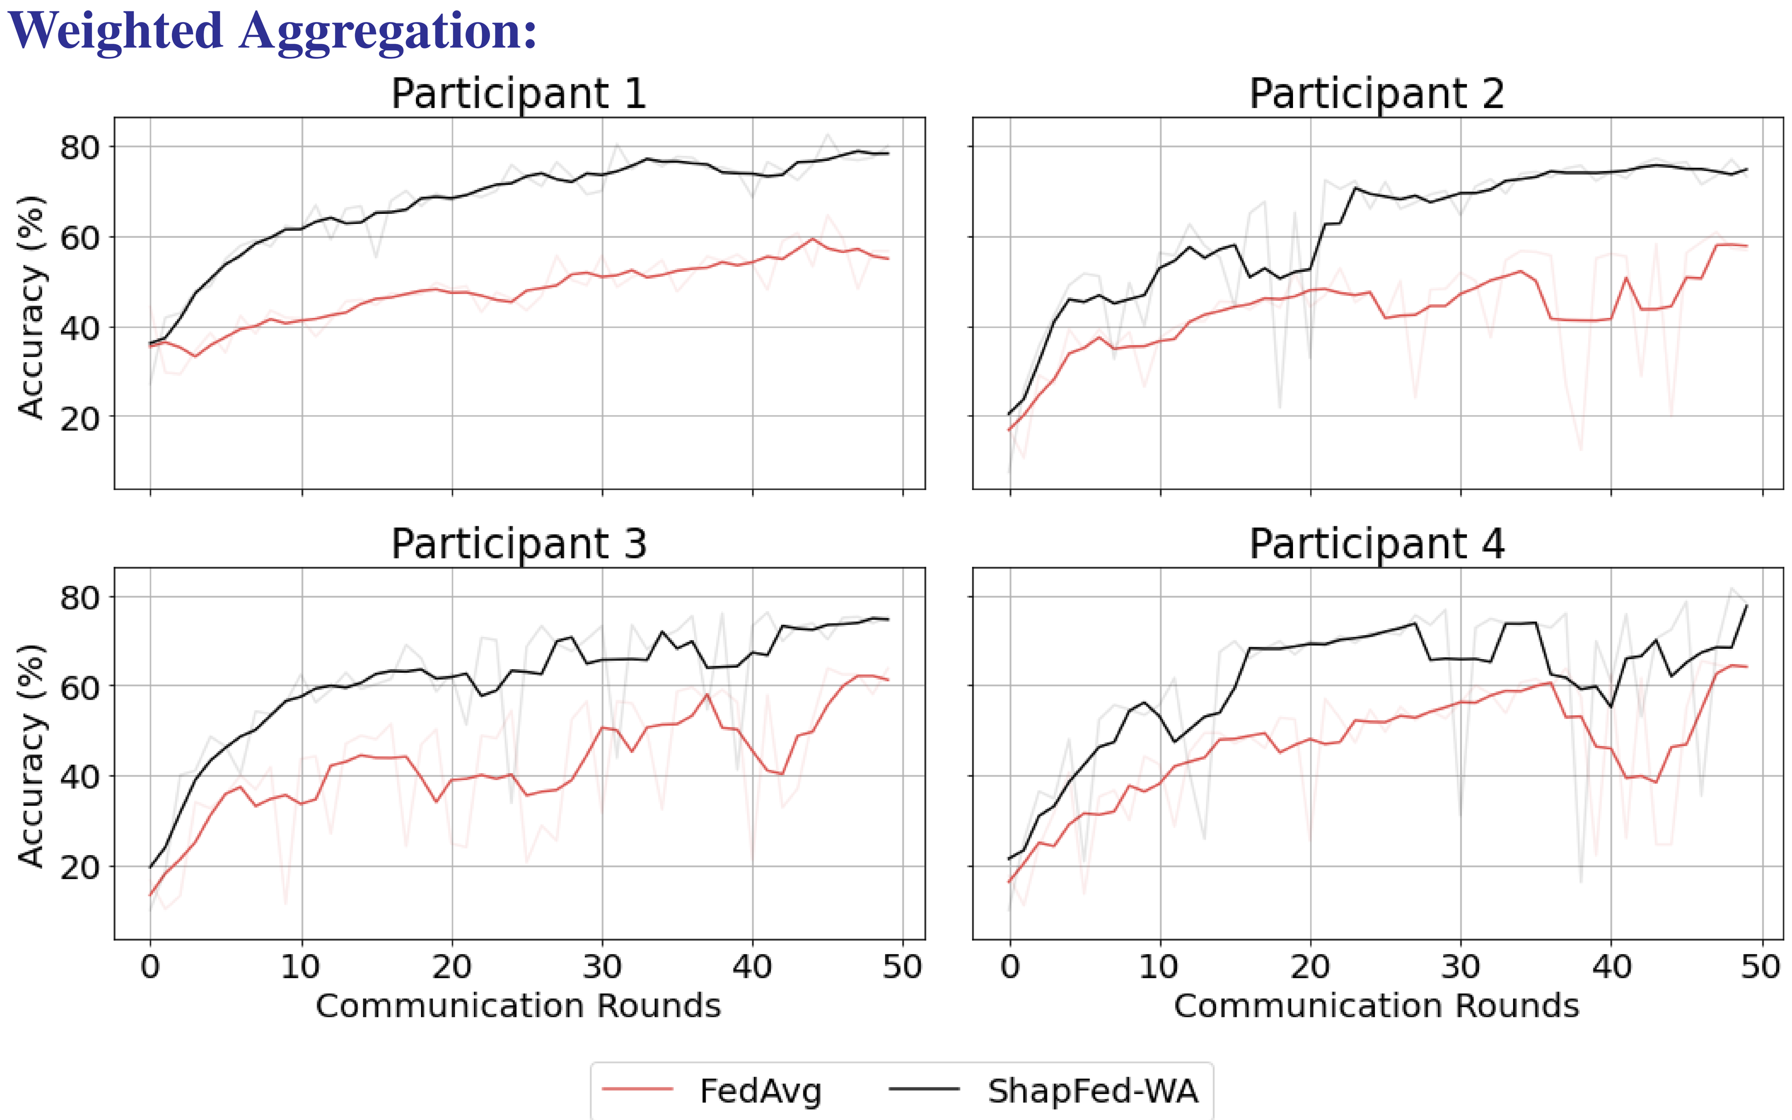

Comparing FedAvg and ShapFed-WA on CIFAR10 under an imbalanced split scenario $(0.7, 1)$ with 4 participants.

Contribution Assessment:

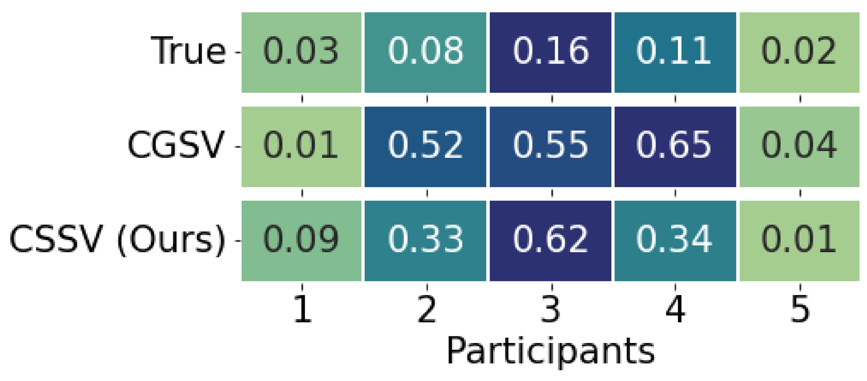

Comparison of our proposed contribution assessment algorithm with CGSV and true Shapley value computations using **ResNet-34 architecture** on **Chest X-Ray** dataset.

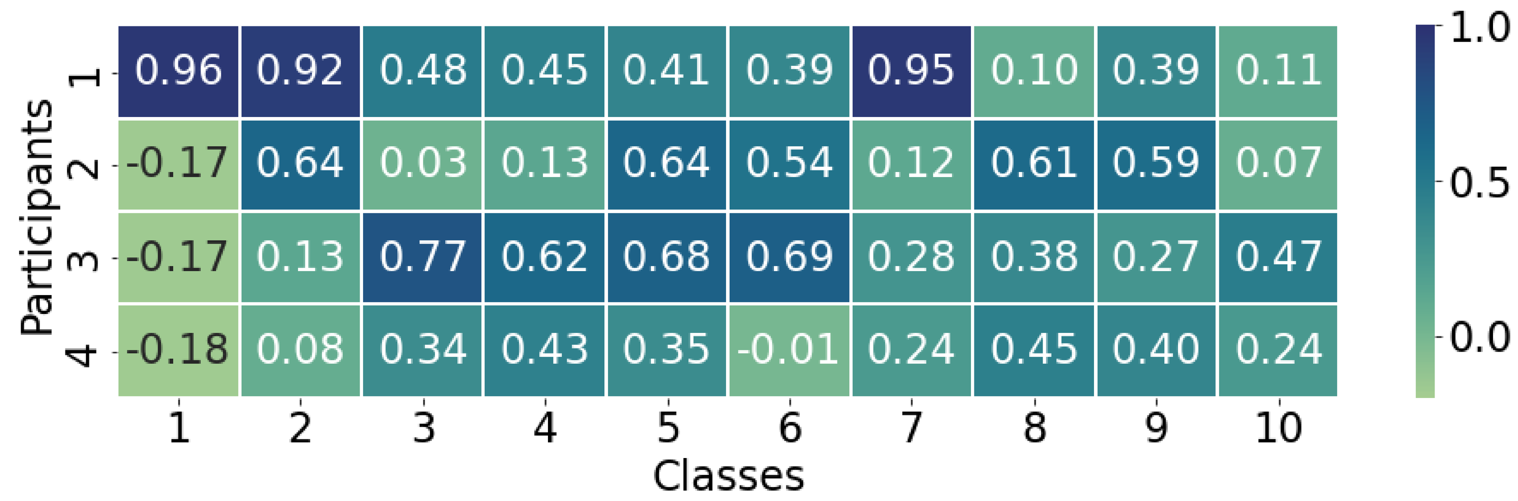

Heatmap visualization of CSSVs for heterogeneous setting evaluated on CIFAR-10 dataset.

Personalization (Fed-ISIC2019):

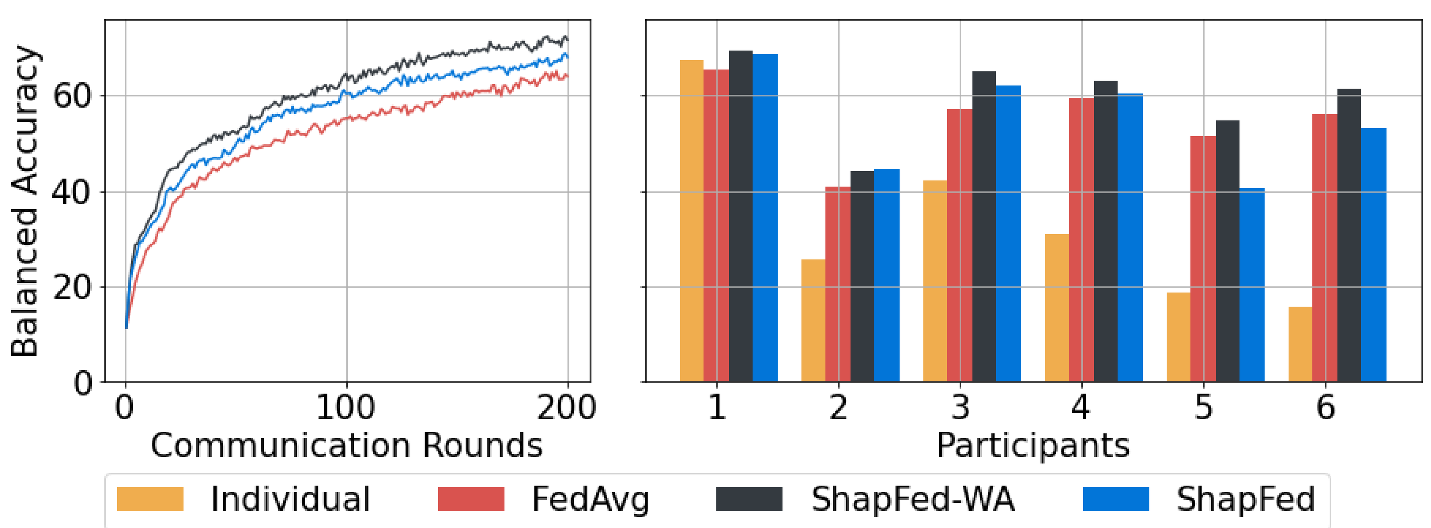

(Left) The balanced accuracy of our methods (ShapFed-WA & ShapFed) vs FedAvg. (Right) Per-participant accuracy on Fed-ISIC2019 dataset.

Setting	P_1	P_2	P_3	P_4	P_5	P_6	Corr.
Individual	67.2	25.7	42.3	31.0	18.5	15.6	—
FedAvg	65.4	40.9	57.2	59.3	51.5	56.2	0.63
ShapFed-WA	69.3	44.3	65.0	63.1	54.8	61.2	0.62
ShapFed	68.5	44.4	61.9	60.4	40.6	53.2	**0.84**

Table for Fed-ISIC2019 experiment results.

Project website.