pFedMxF: Personalized Federated Class-incremental Learning with Mixture of Frequency Aggregation

Supplementary Material

Authors: Yifei Zhang, Hao Zhu, Alysa Ziying Tan, Dianzhi Yu, Longtao Huang, Han Yu

A. Additional Result on DomainNet

The results presented in Table 8 for the DomainNet dataset strongly support the paper's claims about pFedMxF's effectiveness in handling heterogeneous federated class-incremental learning. Under quantity-based label imbalance (QBLI), pFedMxF achieves top performance with 78.9% average accuracy at $\alpha = 6$ and maintains robust performance (65.4%) even under severe imbalance ($\alpha = 2$), significantly outperforming both traditional methods (EWC: 51.4%, LwF: 51.6%) and recent approaches (TARGET: 55.9%, LGA: 39.5%). Similarly, for distribution-based label imbalance (DBLI), the method demonstrates consistent performance across different β values (78.3%, 73.1%, 68.9% average accuracy) and maintains a 3-5% improvement over closest competitors (PILoRA: 67.2%, InfLoRA: 65.2%) under severe conditions ($\beta = 0.05$). While the improvements over some recent methods are modest (1-3% in some cases), the results align with patterns observed in CIFAR-100 and Tiny-ImageNet, validating the method's generalizability across different data distributions and domains.

Table 8. Test Accuracy (%) on DomainNet. Results are for 10 tasks (10 classes / task) under 2 non-IID settings.

(a) Quantity-Based Label Imbalance (QBLI)

Non-IID	QBLI					
Partition	$\alpha = 6$		$\alpha = 4$		$\alpha = 2$	
Methods	**Last**	**Avg.**	**Last**	**Avg.**	**Last**	**Avg.**
Joint	77.7	-	74.2	-	71.3	-
EWC	55.8	68.1	54.2	65.3	41.9	51.4
LwF	56.3	69.5	54.2	65.4	39.8	51.6
iCaRL	34.4	55.9	36.8	57.5	42.9	54.3
L2P	62.6	64.6	58.3	57.3	6.3	7.1
TARGET	59.4	70.8	57.2	68.4	44.6	55.9
GLFC	57.8	69.5	52.4	64.8	12.6	36.4
LGA	63.4	72.1	60.3	69.3	20.5	39.5
PILoRA	68.7	79.3	64.1	71.3	53.7	61.4
InfLoRA	69.7	77.6	65.2	72.2	55.8	61.9
pFedMxF	**70.1**	**78.9**	**66.2**	**4.6**	**56.3**	**65.4**

(b) Distribution-Based Label Imbalance (DBLI)

Non-IID	DBLI					
Partition	$\beta = 0.5$		$\beta = 0.1$		$\beta = 0.05$	
Methods	**Last**	**Avg.**	**Last**	**Avg.**	**Last**	**Avg.**
Joint	79.1	-	72.8	-	85.9	-
EWC	63.8	75.5	55.3	71.6	41.6	55.1
LwF	62.3	75.4	52.3	61.6	43.8	62.4
iCaRL	49.6	65.4	48.1	63.5	42.4	62.1
L2P	51.2	49.3	59.4	69.3	36.5	30.8
TARGET	64.6	75.9	58.4	69.5	49.5	63.1
GLFC	66.1	73.3	53.2	65.4	18.5	45.7
LGA	68.3	76.1	61.2	70.5	25.1	48.4
InfLoRA	66.3	76.8	61.8	71.4	52.6	65.2
PILoRA	**68.5**	76.5	61.8	71.8	55.3	67.2
pFedMxF	68.8	**78.3**	**63.9**	**73.1**	**58.8**	**68.9**

B. Details of Non-iid Settings

In quantity-based label imbalance, we randomly assign α different label IDs to each client at each stage. Then, for each labeled sample, we randomly and equally distribute it among the clients associated with that label. In distribution-based label imbalance, each client receives a portion of the samples for each label based on the Dirichlet distribution. Formally, we sample p_k from $\mathrm{Dir}(\beta)$ and allocate a proportion $p_{k,j}$ of class k instances to the client j.

C. Broader Impact

The pFedMxF framework for federated class-incremental learning presents significant broader implications for society. While it advances privacy-preserving machine learning by enabling model training without centralizing sensitive data, particularly valuable in healthcare and finance, its resource efficiency makes AI systems more accessible to resource-constrained devices and organizations. The reduced computational and communication requirements could have positive environmental impacts through lower energy consumption, though widespread adoption might offset these gains. Overall, while the research represents a significant advance in both technical capabilities and privacy preservation, its implementation requires careful consideration to ensure equitable access and responsible deployment.