

Adaptive Client Sampling by Update Norm in Federated Vision-Language Model Convergence

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 6 peer-reviewed papers addressing the following research question: How does the use of adaptive client sampling based on update norm magnitude in federated learning impact the convergence speed and accuracy of multimodal vision-language models on the VQA-v2. It is well understood that client-master communication can be a primary bottleneck in Federated Learning. In this work, we address this issue with a novel client subsampling scheme, where we restrict the number of clients allowed to communicate their updates back to the master. 17 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Optimal Client Sampling for Federated Learning. Research question: How does the use of adaptive client sampling based on update norm magnitude in federated learning impact the convergence speed and accuracy of multimodal vision-language models on the VQA-v2 benchmark compared to uniform client sampling?.

2 Methodology

Systematic literature search across multiple databases yielded 6 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.8/10.

3 Results

6 papers retrieved. 17 claims extracted; 0 independently verified. Quality review score: 3.8/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The optimal client sampling method is evaluated on standard federated datasets from LEAF (Caldas et al., 2018).	×	0.08
Dataset 1 contains 32,906 training images in total.	×	0.01
In Dataset 1, 10% of clients hold 82% of the training images.	×	0.02
Dataset 2 contains 29,906 training images in total.	×	0.01
In Dataset 2, 20% of clients hold 90% of the training images.	×	0.02
Dataset 3 contains 27,599 training images in total.	×	0.01
In Dataset 3, 50% of clients hold 98% of the training images.	×	0.02
The method is compared with full participation and uniform sampling baselines.	×	0.04
The experiments are conducted using TensorFlow Federated (TFF).	×	0.03
The validation accuracy and training loss are reported as functions of communication rounds and bits communicated.	×	0.03
The results are averaged over 5 independent runs for each method.	×	0.06
The same random seed is used for all methods in a single run.	×	0.03
The code and datasets are available at https://github.com/SamuelHorvath/FL-optimal-client-sampling .	×	0.08
The Federated EMNIST (FEMNIST) dataset is used for image classification.	×	0.04
The training set of FEMNIST is modified to create three unbalanced training sets.	×	0.01
The same CNN model as in (McMahan et al., 2017) is used for the experiments.	×	0.03
The validation set consists of 40,832 images.	×	0.02

References

- <http://arxiv.org/abs/2506.02887v2>

- <http://arxiv.org/abs/2010.13723v3>
- <http://arxiv.org/abs/2207.02337v1>