Published: 12 March 2025
Accepted: 19 February 2025
Revised: 19 February 2025
Received: 19 February 2025

INTRODUCTION

Introduction to the Special Issue on Performance Evaluation of Federated Learning Systems Part 1

CARLEE JOE-WONG, Carnegie Mellon University, Pittsburgh, PA, United States

LILI SU, Northeastern University, Boston, MA, United States

ACM Transactions on Modeling and Performance Evaluation of Computing Systems, Volume 10, Issue 1 (March 2025)
https://doi.org/10.1145/3720399
EISSN: 2376-3647

Introduction to the Special Issue on Performance Evaluation of Federated Learning Systems Part 1

Federated learning has recently emerged as a popular approach to training machine learning models on data that is scattered across multiple heterogeneous devices, often referred to as "clients" in a federated learning system. These clients iteratively compute updates to the machine learning models on their local datasets. These updates are periodically aggregated across clients, typically but not always with the help of a parameter server. The aggregated model then serves as the starting point for new rounds of client updates. In many real-world applications such as connected-and-autonomous vehicles (CAVs), the underlying distributed/decentralized systems on which federated learning algorithms are executing suffer a wide degree of heterogeneity including but not limited to data distributions, computation speeds, and external local environments. Moreover, the clients are often resource-constrained edge or end devices and may compete for common resources such as communication bandwidth. Such heterogeneity raises significant research questions on how these systems will perform under different variants of federated learning algorithms.

This special issue presents ten papers that evaluate the performance, broadly understood, of federated learning systems. The selection process for the special issue was competitive, with an acceptance rate around 30%, and ultimately many worthy submissions could not be included. Each paper published in this special issue went through at least one round of revision to address comments from the reviewers. Given the volume of submissions and selected papers, we decided to split the special issue into two parts, which respectively focus on the performance implications of resource constraints and client data heterogeneity in federated learning.

In Part 1 of this special issue, we present five papers that commonly address the *resource constraints* in federated learning, which can range from constraints on the communication resources between the clients and the parameter server to memory or computation constraints on individual clients. The impacts of these constraints on the performance–including criteria such as model accuracy, resource efficiency or cost, privacy of client data, and client fairness–of federated learning systems is not yet well understood in existing literature, and the selected papers in Part 1 represent important steps forward in quantifying such performance effects.

The first two papers, "*Clipper: Online Joint Client Sampling and Power Allocation for Wireless Federated Learning*" and "*A Green Multi-Attribute Client Selection for Over-the-Air Federated Learning: A Grey-Wolf-Optimizer Approach*" address the specific scenarios in which federated learning clients communicate with a parameter server over a wireless network. In such scenarios, the clients share wireless network resources, which are generally limited and vary in quality due to heterogeneous and dynamic wireless signal strengths at different clients. "*Clipper: Online Joint Client Sampling and Power Allocation for Wireless Federated Learning*" addresses the challenge of allocating limited

ACM Reference Format:
Carlee Joe-Wong and Lili Su. 2025. Introduction to the Special Issue on Performance Evaluation of Federated Learning Systems Part 1. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 10, 1, Article 1 (March 2025), 2 pages.
https://doi.org/10.1145/3720399

ACM 2376-3639/2025/03-ART1
https://doi.org/10.1145/3720399

wireless resources to clients so as to optimize the overall performance of the federated training process. In doing so, the authors develop a novel analysis that quantifies the accuracy of the trained federated learning model in terms of the wireless resource allocation. "*A Green Multi-Attribute Client Selection for Over-the-Air Federated Learning: A Grey-Wolf-Optimizer Approach*" considers a specific wireless federated learning structure in which client updates are aggregated "over-the-air" while being sent to the parameter server. Such an aggregation approach can improve federated learning's privacy guarantees, but it can also introduce significant deployment complexity. This work develops a framework for controlling the number of clients that participate in each federated learning round, thus limiting over-the-air complexity while balancing other performance criteria.

Limited client resources can also impact federated learning performance in practice, particularly when these resources are heterogeneous; that is, different clients have varying computational, memory, and communication resources available to utilize in the federated learning training. "*Review and Comparative Evaluation of Resource-Adaptive Collaborative Training for Heterogeneous Edge Devices*" presents a detailed survey and evaluation of multiple federated learning variants in the presence of such client heterogeneity, including novel empirical findings from a real testbed of heterogeneous devices. "*Hardware-Sensitive Fairness in Heterogeneous Federated Learning*" focuses on the fairness implications of heterogeneous clients in emerging settings where less-resourced federated learning clients might use smaller model architectures. This setting thus raises fairness concerns due to the disparities in model complexity between clients, and the authors propose a novel federated learning framework that trades off client performance and fairness in such heterogeneous settings. Finally, "*SAFE: Secure Aggregation with Failover and Encryption*" focuses on federated learning's privacy guarantees. Introducing traditional secure aggregation into federated learning often comes with high computing and communication overhead, which may be prohibitive for some clients. The authors thus propose an alternative secure aggregation framework that relies on "chaining" clients for secure model aggregation, which is shown to scale better than traditional approaches.

We would like to conclude by thanking the TOMPECS Editor-in-Chief, Leana Golubchik, for inviting us to organize this special issue; Gita Delsing, the TOMPECS journal administrator; all of the submission authors, for their fine contributions; and finally all of the reviewers, whose comments were invaluable to the review and revision process. We hope that you enjoy reading Part 1 of our special issue!

Carlee Joe-Wong
Carnegie Mellon University, Pittsburgh, United States
email: cjoewong@andrew.cmu.edu

Lili Su
Northeastern University, Boston, United States
email: l.su@northeastern.edu

Received 19 February 2025; revised 19 February 2025; accepted 19 February 2025