Qeios

Peer Review

Review of: "One Communication Round is All It Needs for Federated Fine-Tuning Foundation Models"

Pian Qi[1]

1. University of Naples Federico II, Italy

This work provides both theoretical and empirical evidence that a single round of aggregation can yield a global model performance comparable to multiple rounds of aggregation. The experiments are well-conducted, and the results are thoroughly analyzed. However, a few aspects could be further explored to strengthen the study:

1. Comparison with Other PEFT Methods

- In the federated learning experiments, LoRA was used as the primary parameter-efficient fine-tuning method. However, how does the proposed approach compare against other PEFT techniques such as Adapters, Prefix-Tuning, or BitFit? Including these methods would provide a more comprehensive benchmark.

2. Generalization to Domain-Specific Tasks

- The dataset used for fine-tuning appears to be relatively general, likely covering everyday scenarios. Would the results hold for domain-specific downstream tasks in areas such as healthcare (e.g., MIMIC-III), industrial automation, or transportation (e.g., Waymo Open Dataset)? Since FL is often deployed in specialized fields with unique data distributions, evaluating performance on such datasets would enhance the study's practical relevance.

3. Distributed Simulation of FL

- Did the authors consider performing distributed FL simulations using platforms like FLOWER or FedML? Such an environment could better simulate the challenges of real-world FL, including heterogeneous client data, varying communication constraints, and model upload/download latencies. Conducting experiments in a more realistic distributed setting would further validate the findings and strengthen their applicability.

 doi.org/10.32388/NDLKHN

Declarations

Potential competing interests: No potential competing interests to declare.