Comparison of FedKRSO and Standard LoRA FL on SuperGLUE WSC and RTE

Assignee Research

June 12, 2026

Abstract

Fine-tuning large language models requires high computational and memory resources, and is therefore associated with significant costs. When training on federated datasets, an increased communication effort is also needed. For this reason, parameter-efficient methods (PEFT) are becoming increasingly important. In this context, very good results have already been achieved by fine-tuning with low-rank adaptation methods (LoRA). The application of LoRA methods in Federated Learning, and especially the aggregation of adaptation matrices, is a current research field. In this article, we propose a n

1 Introduction

This paper examines: Aggregating Low Rank Adapters in Federated Fine-tuning. Research question: How does FedKRSO compare to standard LoRA FL in terms of convergence speed and final accuracy on the WSC and RTE subsets of SuperGLUE?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 8.5/10.

3 Results

11 papers retrieved. 14 claims extracted; 14 independently verified. Quality review score: 8.5/10.

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
FFA and FRA-LoRA are the only options for privacy-preserving training among the methods discussed.	✓	0.18
On the SST2 dataset with imbalanced classes, FRA-LoRA performs well in the first 15-20 communication rounds before accur	✓	0.20
For the SST2 dataset, the Centralised method achieved an evaluation set accuracy of 94.7248.	✓	0.18
For the SST2 dataset, FedAvg achieved an evaluation set accuracy of 93.1193, outperforming FRA (92.7752) and FFA (92.545	✓	0.22
For the MNLI dataset, the Centralised method achieved an evaluation set accuracy of 86.0316.	✓	0.16
For the MNLI dataset, FFA achieved an accuracy of 83.5354, which is lower than FedAvg (85.7769) and FRA-LoRA (85.5018).	✓	0.17
In experiments on the SST2 dataset, a dropout factor of 0.2, weight decay of 0.005, and a learning rate of 0.0001 were u	✓	0.19
FRA-LoRA starts overfitting earlier than FFA-LoRA on smaller datasets, potentially due to having double the number of pa	✓	0.15
On the imbalanced MNLI dataset, vanilla Federated Averaging and FRA-LoRA perform equally well.	✓	0.25
Only FRA-LoRA is compatible with additional noise among the compared methods.	✓	0.16
The error in FedAvg arises from inequality (1) and is calculated as errFedAvg = 0.25 $\cdot$ (A1 - A2) $\cdot$ (B2 - B1).	✓	0.20
FRA-LoRA errors are an order of magnitude lower in absolute terms compared to FedAvg errors in the MNLI imbalanced examp	✓	0.30
The EU AI Act states energy sustainability as a requirement for AI systems.	✓	0.18
LoRA freezes actual weight matrices and trains only two low-rank adapter matrices whose product approximates the gradien	✓	0.17

References

- http://arxiv.org/abs/2306.06371v1
- http://arxiv.org/abs/2501.06332v1
- http://arxiv.org/abs/2207.08179v1