

# How does varying LoRA rank impact cross-domain transfer accuracy on GLUE semantic textual similarity tasks und

Assignee Research

June 10, 2026

## Abstract

Foundation models (FMs) achieve strong performance across diverse tasks with task-specific fine-tuning, yet full parameter fine-tuning is often computationally prohibitive for large models. Parameter-efficient fine-tuning (PEFT) methods like Low-Rank Adaptation (LoRA) reduce this cost by introducing low-rank matrices for tuning fewer parameters. While LoRA allows for efficient fine-tuning, it requires significant data for adaptation, making Federated Learning (FL) an appealing solution due to its privacy-preserving collaborative framework. However, combining LoRA with FL introduces two key cha

## 1 Introduction

This paper examines: LoRA-FAIR: Federated LoRA Fine-Tuning with Aggregation and Initialization Refinement. Research question: How does varying LoRA rank impact cross-domain transfer accuracy on GLUE semantic textual similarity tasks under differential privacy constraints?.

## 2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.8/10.

## 3 Results

15 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.8/10.

## 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.

## References

- <http://arxiv.org/abs/2412.09827v1>
- <http://arxiv.org/abs/2410.09437v3>
- <http://arxiv.org/abs/2411.14961v3>