

CausalMixFT Privacy-Utility Trade-offs vs. Differentially Private Fine-Tuning Methods

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the privacy-utility trade-off of CausalMixFT compare to other differentially private fine-tuning methods (e.g., DP-AdamW, DP-LAMB) when evaluated on standard NLP benchmarks (e.g., GLUE). 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 7.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Differentially Private Fine-tuning of Language Models. Research question: How does the privacy-utility trade-off of CausalMixFT compare to other differentially private fine-tuning methods (e.g., DP-AdamW, DP-LAMB) when evaluated on standard NLP benchmarks (e.g., GLUE, SuperGLUE) at varying privacy budgets (ϵ values)?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 7.5/10.

3 Results

14 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 7.5/10.

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.

References

- <http://arxiv.org/abs/2511.07843v1>
- <http://arxiv.org/abs/2511.20509v2>
- <http://arxiv.org/abs/2110.06500v2>