LongLoRA Shifted Sparse Attention and Its Effects on HumanEval Pass@1 Performance

Assignee Research

June 1, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: What is the impact of LongLoRA's shifted sparse attention mechanism on pass@1 scores in HumanEval compared to full fine-tuning and adapter-based methods. We present LongLoRA, an efficient fine-tuning approach that extends the context sizes of pre-trained large language models (LLMs), with limited computation cost. Typically, training LLMs with long context sizes is computationally expensive, requiring extensive training hours and. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LongLoRA: Efficient Fine-tuning of Long-Context Large Language Models. Research question: What is the impact of LongLoRA's shifted sparse attention mechanism on pass@1 scores in HumanEval compared to full fine-tuning and adapter-based methods?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.0/10.

3 Results

8 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.0/10.

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/2512.07011v1
- http://arxiv.org/abs/2303.16199v3
- http://arxiv.org/abs/2309.12307v3