Block-Sparse FlashAttention Performance and Efficiency at 256k-Token Context Lengths

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: Can Block-Sparse FlashAttention maintain its retrieval performance on cross-lingual retrieval tasks like MTOP when context lengths exceed 256k tokens, and how does its efficiency scale compared to. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Block Sparse Flash Attention. Research question: Can Block-Sparse FlashAttention maintain its retrieval performance on cross-lingual retrieval tasks like MTOP when context lengths exceed 256k tokens, and how does its efficiency scale compared to dense attention?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.2/10.

3 Results

13 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.2/10.

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/2605.11733v1
- http://arxiv.org/abs/2512.07011v1
- http://arxiv.org/abs/2601.15305v1