

LongNet and H3 Perplexity Degradation on PG-19 Beyond 64k Tokens

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: What is the perplexity degradation of LongNet versus H3 on the PG-19 benchmark when evaluated at sequence lengths exceeding 64k tokens. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LongNet: Scaling Transformers to 1,000,000,000 Tokens. Research question: What is the perplexity degradation of LongNet versus H3 on the PG-19 benchmark when evaluated at sequence lengths exceeding 64k tokens?.

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 3.2/10.

3 Results

11 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.2/10.

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/2502.17129v2>
- <http://arxiv.org/abs/2307.02486v2>
- <http://arxiv.org/abs/2505.09561v2>