

Robustness of Llama-3-8B-128K Qwen-8B and Mistral-8B to Irrelevant Context Noise at 128K Scale

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How do Llama-3-8B-128K, Qwen-8B, and Mistral-8B differ in robustness to irrelevant context noise within the BABILong dataset as the total sequence length scales to 128K. We study the continual pretraining recipe for scaling language models' context lengths to 128K, with a focus on data engineering. We hypothesize that long context modeling, in particular the ability to utilize information at arbitrary input locations, is a capability. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Data Engineering for Scaling Language Models to 128K Context. Research question: How do Llama-3-8B-128K, Qwen-8B, and Mistral-8B differ in robustness to irrelevant context noise within the BABILong dataset as the total sequence length scales to 128K?.

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 6.5/10.

3 Results

14 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.5/10.

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/2511.05080v4>
- <http://arxiv.org/abs/1404.0103v1>
- <http://arxiv.org/abs/2402.10171v1>