

# Bilingual Context Mixing Effects on Attention Sink Mitigation in Llama-3-8B with FlowKV and DynamicRoPE

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: What is the impact of bilingual context mixing (English-Chinese) on the attention sink mitigation effectiveness of FlowKV versus DynamicRoPE in Llama-3-8B when evaluated on LongBench Pro. 15 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: LongBench Pro: A More Realistic and Comprehensive Bilingual Long-Context Evaluation Benchmark. Research question: What is the impact of bilingual context mixing (English-Chinese) on the attention sink mitigation effectiveness of FlowKV versus DynamicRoPE in Llama-3-8B when evaluated on LongBench Pro?.

## 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 3.3/10.

## 3 Results

8 papers retrieved. 15 claims extracted; 0 independently verified. Quality review score: 3.3/10.

## 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.

## 5 Extracted Claims

Claim	Verified	Confidence
LongBench Pro includes 11 primary task categories.	×	0.07
LongBench Pro introduces the context requirement dimension with two categories: Full and Partial.	×	0.08
LongBench Pro crosses 11 primary categories with context requirements to design 25 secondary categories.	×	0.10
LongBench Pro documents are collected from diverse domains and formats, including news, medicine, science, literature, l	×	0.04
LongBench Pro documents are balanced across single-document and multi-document settings, as well as English and Chinese.	×	0.08
LongBench Pro documents are assigned to six target length buckets: 8k, 16k, 32k, 64k, 128k, and 256k tokens.	×	0.07
LongBench Pro documents undergo a compliance review by human annotators to exclude privacy-sensitive, copyrighted, or no	×	0.02
LongBench Pro uses two types of prompts: non-thinking and thinking prompts.	×	0.04
LongBench Pro collects predictions from five advanced models for each sample during the answer review process.	×	0.02
LongBench Pro samples are reviewed by two annotators independently, and samples with detected issues are evaluated by an	×	0.09
Gemini-2.5-Pro has an overall score of 55.92 in the LongBench Pro benchmark.	×	0.06
Gemini-2.5-Flash has an overall score of 36.14 in the LongBench Pro benchmark.	×	0.04
Gemma-3-27B-It has an overall score of 32.16 in the LongBench Pro benchmark.	×	0.04
Gemma-3-12B-It has an overall score of 21.76 in the LongBench Pro benchmark.	×	0.03
Gemma-3-4B-It has an overall score of 20.89 in the LongBench Pro benchmark.	×	0.04

## References

- <http://arxiv.org/abs/2308.14508v2>
- <http://arxiv.org/abs/2601.02872v1>
- <http://arxiv.org/abs/2601.15305v1>