

# Context-Length Scaling in CAKE Code Generation Accuracy on HumanEval-Multi

Assignee Research

June 7, 2026

## Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: What is the degradation in CAKE’s code generation accuracy on HumanEval-Multi when context length exceeds 32k tokens relative to full attention baselines. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.8/10. 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: What is the degradation in CAKE’s code generation accuracy on HumanEval-Multi when context length exceeds 32k tokens relative to full attention baselines?.

## 2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.8/10.

## 3 Results

10 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.8/10.

## 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/2602.02159v1>
- <http://arxiv.org/abs/2601.15305v1>