

KAN and MLP Inference Latency on the HellaSwag Reasoning Benchmark

Assignee Research

June 4, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the inference latency of KANs compare to traditional MLPs when evaluated on the HellaSwag reasoning benchmark for language models. 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: VIKIN: A Reconfigurable Accelerator for KANs and MLPs with Two-Stage Sparsity Support. Research question: How does the inference latency of KANs compare to traditional MLPs when evaluated on the HellaSwag reasoning benchmark for language models?.

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/2603.01165v1>
- <http://arxiv.org/abs/2405.14399v1>
- <http://arxiv.org/abs/2412.16738v1>