

# Computational Efficiency Trade-offs in Scaling Kimi k1.5 RL Training vs R1-7B Pretraining

Assignee Research

June 6, 2026

## Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: What is the computational efficiency trade-off when scaling Kimi k1.5's RL-based training compared to R1-7B's pretraining, measured by throughput on the GEMM or PopQA benchmarks. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Kimi Linear: An Expressive, Efficient Attention Architecture. Research question: What is the computational efficiency trade-off when scaling Kimi k1.5's RL-based training compared to R1-7B's pretraining, measured by throughput on the GEMM or PopQA benchmarks?.

## 2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.2/10.

## 3 Results

4 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 2.2/10.

## 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/2510.26692v2>
- <http://arxiv.org/abs/2507.20534v2>
- <http://arxiv.org/abs/2504.07491v3>