

Llama-3.1-8B Performance on LiveCodeBench Under Low-Resource Constraints

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How does Llama-3.1-8B's performance on LiveCodeBench compare to smaller or similarly sized language models when evaluated under low-resource conditions or limited inference budgets. Large Language Models achieve remarkable performance but incur substantial computational costs unsuitable for resource-constrained deployments. This paper presents the first comprehensive task-specific efficiency analysis comparing 16 language models across five diverse NLP. 3 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Task-Specific Efficiency Analysis: When Small Language Models Outperform Large Language Models. Research question: How does Llama-3.1-8B's performance on LiveCodeBench compare to smaller or similarly sized language models when evaluated under low-resource conditions or limited inference budgets?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.5/10.

3 Results

12 papers retrieved. 3 claims extracted; 0 independently verified. Quality review score: 2.5/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
The Qwen2.5-0.5B model achieved an accuracy of 0.3768 on the GSM8K benchmark.	×	0.03
The Qwen2.5-72B model has a latency of 0.15 ms/token on the GSM8K benchmark.	×	0.02
The Qwen2.5-0.5B model achieved an accuracy of 0.39 on the HellaSwag benchmark.	×	0.03

References

- <http://arxiv.org/abs/2603.21389v1>
- <http://arxiv.org/abs/2504.04596v1>
- <http://arxiv.org/abs/2505.21514v1>