

Model Size Effects on Self-Repair Efficiency in Llama-2 Code Generation

Assignee Research

June 3, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: To what extent does model size (e.g., 7B vs. 13B vs. 70B) impact the efficiency of self-repair in Llama-2 models, evaluated by the trade-off between pass@1 accuracy and inference latency in code. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Is Self-Repair a Silver Bullet for Code Generation?. Research question: To what extent does model size (e.g., 7B vs. 13B vs. 70B) impact the efficiency of self-repair in Llama-2 models, evaluated by the trade-off between pass@1 accuracy and inference latency in code generation tasks?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.8/10.

3 Results

11 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.8/10.

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/2306.09896v5>
- <http://arxiv.org/abs/2310.20624v2>
- <http://arxiv.org/abs/2604.10508v1>