

Codestral Model Size and Inference Latency in Big-Vul Severity Classification

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: What is the correlation between model size and inference latency for Codestral when performing severity-level classification on C and C++ code in Big-Vul. Context: Traditional software security analysis methods struggle to keep pace with the scale and complexity of modern codebases, requiring intelligent automation to detect, assess, and remediate vulnerabilities more efficiently and accurately. Objective: This paper explores the. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: LLMs in Code Vulnerability Analysis: A Proof of Concept. Research question: What is the correlation between model size and inference latency for Codestral when performing severity-level classification on C and C++ code in Big-Vul?.

2 Methodology

Systematic literature search across multiple databases yielded 16 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.0/10.

3 Results

16 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.0/10.

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/2506.11022v2>
- <http://arxiv.org/abs/2509.17337v1>
- <http://arxiv.org/abs/2601.08691v1>