

Retrieval-Augmented Generation Impact on False Positives in Malicious Python Code Detection

Assignee Research

June 4, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: How does Retrieval-Augmented Generation affect the false positive rates of Llama 3.1 compared to Mistral 7B when classifying obfuscated Python code in malicious package detection. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Detecting Malicious Source Code in PyPI Packages with LLMs: Does RAG Come in Handy?. Research question: How does Retrieval-Augmented Generation affect the false positive rates of Llama 3.1 compared to Mistral 7B when classifying obfuscated Python code in malicious package detection?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.3/10.

3 Results

10 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 4.3/10.

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/2310.06825v1>
- <http://arxiv.org/abs/2402.12317v2>
- <http://arxiv.org/abs/2504.13769v1>