

Mixed C/C++ and Python Training Data Effects on Security Vulnerability Detection in Codestral-7B and Llama3-70B

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: What is the impact of mixing C/C++ and Python code samples in the training set on the F1 score for security vulnerability detection tasks in Codestral-7B versus Llama3-70B. Large Language Models (LLMs) have demonstrated significant capabilities in understanding and analyzing code for security vulnerabilities, such as Common Weakness Enumerations (CWEs). However, their reliance on cloud infrastructure and substantial computational requirements pose. 13 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Case Study: Fine-tuning Small Language Models for Accurate and Private CWE Detection in Python Code. Research question: What is the impact of mixing C/C++ and Python code samples in the training set on the F1 score for security vulnerability detection tasks in Codestral-7B versus Llama3-70B?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

14 papers retrieved. 13 claims extracted; 2 independently verified. Quality review score: 4.5/10.

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
Farasat & Posegga [15] achieved an average accuracy of 98.6%, F1 score of 94.7%, precision of 96.2%, recall of 93.3%, an	×	0.09
Bagheri et al. [31] reported an F1 score of 99% using a hybrid ML model (Self-attention + CNN - Conformer).	×	0.05
Singh et al. [32] achieved an accuracy of 0.66, precision of 0.65, recall of 0.66, and F1 score of 0.64 for CWE predicti	×	0.04
Dozono et al. [6] reported Python F1 scores for various LLMs: GPT-4o=0.80, GPT-4T=0.76, Gemini 1.5 Pro=0.75, CodeLlama-7	×	0.01
Steenhoek et al. [34] evaluated LLMs on C/C++ and found a low balanced accuracy of 54.5%.	×	0.04
Shestov et al. [35] achieved an F1 score of 0.86 for binary classification using WizardCoder for JAVA CWE detection.	×	0.05
Li et al. [36] reported a 5-6% improvement over the base model using LoRa and IA3 fine-tuning approach for LLMs.	×	0.06
Jiang et al. [37] achieved the best F1 score of 87% using the Llama 2-7b model with LoRa fine-tuning.	×	0.06
The un-tuned codegen-mono model failed to detect a single CWE within any of the 100 code snippets presented in the basel	×	0.12
After instruction-following fine-tuning, the codegen-mono model achieved an accuracy of 99%, precision of $\approx 98.08\%$, recal	✓	0.23
The methodology involves fine-tuning a Small Language Model (SLM) to detect specific Common Weakness Enumerations (CWEs)	✓	0.21
The target vulnerabilities for the detection model are selected from the MITRE Top 25 Most Dangerous Software Weaknesses	×	0.05
The dataset curation process involves a semi-supervised approach combining automated generation with human oversight usi	×	0.06

References

- <http://arxiv.org/abs/2504.16584v1>
- <http://arxiv.org/abs/2507.16540v1>
- <http://arxiv.org/abs/2408.15301v2>