

Instruction Fine-Tuning Effects on Llama3 Efficiency and Accuracy for Obfuscated Code Detection

Assignee Research

June 4, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does instruction fine-tuning on domain-specific security datasets impact the inference efficiency and detection accuracy trade-off for Llama3-7B and Llama3-70B on obfuscated code samples. 12 claims were extracted from source literature; 1 was 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: How does instruction fine-tuning on domain-specific security datasets impact the inference efficiency and detection accuracy trade-off for Llama3-7B and Llama3-70B on obfuscated code samples?.

2 Methodology

Systematic literature search across multiple databases yielded 13 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

13 papers retrieved. 12 claims extracted; 1 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). | × | 0.03 |
| 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 | × | 0.03 |
| Steenhoek et al. [34] evaluated LLMs on C/C++ and found a low balanced accuracy of 54.5%. | × | 0.01 |
| Shestov et al. [35] achieved an F1 score of 0.86 for binary classification using WizardCoder for JAVA CWE detection. | × | 0.06 |
| Li et al. [36] reported a 5-6% improvement over the base model using LoRa and IA3 fine-tuning approach for LLMs. | × | 0.05 |
| 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 | × | 0.13 |
| After instruction-following fine-tuning, the codegen-mono model achieved an accuracy of 99%, precision of $\approx 98.08\%$, recal | ✓ | 0.22 |
| The MITRE Top 25 Most Dangerous Software Weaknesses list [30] was selected as the target for the detection model. | × | 0.05 |
| Google’s gemini-2.0-flash-thinking-exp-01-21 model was used to generate synthetic data for the dataset. | × | 0.03 |

References

- <http://arxiv.org/abs/2408.07888v2>

- <http://arxiv.org/abs/2312.10793v3>
- <http://arxiv.org/abs/2504.16584v1>