

# Taxonomy-Aligned Vulnerability Fine-Tuning for Zero-Shot Code Repair Generalization

Assignee Research

May 31, 2026

## Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: To what extent does taxonomy-aligned vulnerability fine-tuning improve zero-shot generalization on out-of-distribution code repair benchmarks like QuixBugs versus general code corpora. As Large Language Models (LLMs) become increasingly integrated into secure software development workflows, a critical question remains unanswered: can these models not only detect insecure code but also reliably classify vulnerabilities according to standardized taxonomies? In. 11 claims were extracted from source literature; 5 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.4/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Can Open Large Language Models Catch Vulnerabilities?. Research question: To what extent does taxonomy-aligned vulnerability fine-tuning improve zero-shot generalization on out-of-distribution code repair benchmarks like QuixBugs versus general code corpora?.

## 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 6.4/10.

## 3 Results

14 papers retrieved. 11 claims extracted; 5 independently verified. Quality review score: 6.4/10.

## 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
The study evaluates three LLMs: Llama3, Codestral, and Deepseek R1.	✓	0.15
The evaluation uses a carefully filtered subset of the Big-Vul dataset.	✓	0.17
The dataset subset is annotated with eight representative Common Weakness Enumeration (CWE) categories.	✓	0.18
The study adopts a closed-world classification setup.	×	0.13
The study assesses each model’s performance in identifying the presence of vulnerabilities.	×	0.15
The study assesses each model’s performance in mapping vulnerabilities to the correct CWE label.	×	0.14
The evaluated models demonstrated high detection rates for vulnerabilities.	×	0.10
The evaluated models demonstrated markedly poor classification accuracy for CWE labels.	×	0.13
The models exhibited frequent overgeneralization and misclassification of vulnerabilities.	×	0.12
The study analyzes model-specific biases and common failure modes.	✓	0.17
LLMs are being adopted as learning aids in educational contexts.	✓	0.18

## References

- <https://doi.org/10.1007/s11704-024-40231-1>
- <https://doi.org/10.4230/oasics.icpec.2025.4>
- <https://doi.org/10.1145/3475960.3475985>