

Quantization Effects on False Positive Rates in LLaMA 3.2 and Mistral on BugsInPy

Assignee Research

June 11, 2026

Abstract

Large language models (LLMs) have demonstrated strong performance on a wide range of software engineering tasks, including code generation and analysis. However, most prior work relies on cloud-based models or specialized hardware, limiting practical applicability in privacy-sensitive or resource-constrained environments. In this paper, we present a systematic empirical evaluation of two locally deployed LLMs, LLaMA 3.2 and Mistral, for real-world Python bug detection using the BugsInPy benchmark. We evaluate 349 bugs across 17 projects using a zero-shot prompting approach at the function level.

1 Introduction

This paper examines: An Empirical Evaluation of Locally Deployed LLMs for Bug Detection in Python Code. Research question: How does 4-bit quantization affect the false positive rate of LLaMA 3.2 versus Mistral on the BugsInPy dataset compared to FP16 precision?.

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 7.9/10.

3 Results

13 papers retrieved. 24 claims extracted; 21 independently verified. Quality review score: 7.9/10.

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
Chen et al. introduced Codex, a GPT-based model fine-tuned on publicly available code.	✓	0.25
Codex was evaluated on its ability to generate functional Python programs from natural language descriptions using the H	✓	0.24
Feng et al. proposed CodeBERT, which was trained jointly on natural language and programming language data.	✓	0.28
CodeBERT enables tasks such as code search and documentation generation.	×	0.13
Kang et al. proposed AutoFL, which uses LLMs for fault localization while generating natural language explanations along	✓	0.32
Mhatre et al. evaluated cloud-based models on a benchmark spanning both Python and C++.	✓	0.20
Mhatre et al. found that defect complexity is the primary factor governing detection accuracy.	✓	0.22
Prior studies converge on the observation that LLM performance degrades when bugs require cross-function reasoning.	✓	0.18
Santana et al. measured how well LLMs detect and refactor test smells, reporting variation across model families and tas	✓	0.29
Acharya and Ginde applied instruction-tuned models to convert unstructured bug reports into structured form.	✓	0.23
Acharya and Ginde found that open-weight models can approach proprietary system performance on converting bug reports, t	✓	0.20
Widyasari et al. released BugsInPy, a curated set of real Python bugs with reproducible test cases drawn from well-known	✓	0.32
Aguilar et al. identified reproducibility issues in the BugsInPy dataset and proposed revisions to support more reliable	✓	0.23
Pushkar et al. extended evaluation to multi-vulnerability settings, showing that model performance drops systematically	✓	0.28
Zhang et al. contributed a systematic review of LLM-based automated program repair, categorizing methods and identifying	✓	0.27
The availability of open-weight models such as LLaMA and Mistral has made local deployment a practical option.	✓	0.20
Local deployment of LLMs eliminates cloud dependencies and usage costs.	×	0.12
Prior work has largely evaluated large cloud-hosted models rather than locally deployed ones.	✓	0.16
Locally executed models achieve accuracy between 42% and 45% in bug detection	✓	0.24

References

- <http://arxiv.org/abs/2604.23361v1>
- <http://arxiv.org/abs/2411.10958v7>
- <http://arxiv.org/abs/2605.11396v1>