

How does 4-bit quantization impact the syntactic robustness of LLaMA 3.2 and Mistral when generating structure

Assignee Research

June 10, 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 impact the syntactic robustness of LLaMA 3.2 and Mistral when generating structured bug reports compared to FP16 precision?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.8/10.

3 Results

12 papers retrieved. 10 claims extracted; 4 independently verified. Quality review score: 5.8/10.

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
Locally executed models achieve accuracy between 43% and 45% in bug detection.	✓	0.26
Locally executed models produce a large proportion of partially correct responses that identify problematic code regions	✓	0.32
Performance varies significantly across projects, highlighting the importance of codebase characteristics.	✓	0.23
Local models can identify a meaningful share of bugs, though precise localization remains difficult for locally executed	✓	0.34
Performance degrades when bugs require cross-function reasoning.	×	0.05
Defect complexity is the primary factor governing detection accuracy.	×	0.03
Model performance drops systematically as the number of co-occurring defects increases.	×	0.04
Open-weight models can approach proprietary system performance on converting unstructured bug reports into structured fo	×	0.06
The availability of open-weight models such as LLaMA and Mistral has made local deployment a practical option.	×	0.06
Prior work has largely evaluated large cloud-hosted models.	×	0.08

References

- <http://arxiv.org/abs/2406.04710v2>
- <http://arxiv.org/abs/2605.11396v1>
- <http://arxiv.org/abs/2604.23361v1>