Zero-Shot Anomaly Detection in Tabular Data: Llama 3.1 8B vs. Mistral 7B Performance

Assignee Research

June 5, 2026

Abstract

This report synthesises findings from 8 peer-reviewed papers addressing the following research question: How does the zero-shot anomaly detection precision-recall performance of Llama 3.1 8B compare to Mistral 7B when evaluated on synthetic tabular datasets with varying degrees of feature correlation. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: A Review of Large Language Models for Energy Systems: Applications, Challenges, and Future Prospects. Research question: How does the zero-shot anomaly detection precision-recall performance of Llama 3.1 8B compare to Mistral 7B when evaluated on synthetic tabular datasets with varying degrees of feature correlation?.

2 Methodology

Systematic literature search across multiple databases yielded 8 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

8 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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.

References

- https://doi.org/10.1109/access.2025.3610994
- https://doi.org/10.48550/arxiv.2312.01678
- https://doi.org/10.48550/arxiv.2502.01812