Retrieval-Augmented Generation Performance in Zero-Shot Industrial Anomaly Detection

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: How does retrieval-augmented generation with Mistral 7B compare to Llama3.1-8B in zero-shot anomaly detection accuracy on industrial time-series datasets. 6 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: Zero-Shot Anomaly Detection with Pre-trained Segmentation Models. Research question: How does retrieval-augmented generation with Mistral 7B compare to Llama3.1-8B in zero-shot anomaly detection accuracy on industrial time-series datasets?.

2 Methodology

Systematic literature search across multiple databases yielded 11 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

11 papers retrieved. 6 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.

5 Extracted Claims

Claim	Verified	Confidence
The F1-score at the optimal threshold (F1-max) is used to assess the method, as it is less influenced by class imbalance	×	0.05
The method achieves a new state-of-the-art for sample-wise F1-max (81.5) on the VisA dataset.	×	0.11
The method's pixel-wise performance greatly improves over the baseline on the VisA dataset.	×	0.06
The method combines dichotomous image segmentation with SegmentAnything (SAM) to identify the foreground of the image.	×	0.06
Using dichotomous segmentation alone often overestimates the size of an object to include part of the shadows.	×	0.02
SAM normally produces a separate annotation for the object and its shadow.	×	0.02

References

- http://arxiv.org/abs/2411.15497v3
- http://arxiv.org/abs/2306.09269v1
- http://arxiv.org/abs/2504.08115v1