Multi-Modal Outage Detection via Graph-Enhanced Retrieval

Udaivir Yadav
Microsoft

Francisco Mandujano-Reyes
Microsoft

Youjiang Wu
Microsoft

Abstract

Outage detection in cloud services is challenging because failure signals are scattered across heterogeneous sources: time series telemetry, alerts events from a variety of monitors, customer support requests etc. which are incomplete when analyzed independently. We propose a Graph-enhanced retrieval framework that unifies these multi-modal signals for accurate detection. Time series are encoded as vector embeddings, and a knowledge graph captures relationships among services, monitors, root causes, and historical incident cascades. An iterative retrieval mechanism fuses local and global patterns, achieving accuracy and improving recall over time-series baselines and yielding actionable root-causes. Our system's reasoning supports earlier detection and enables lower mitigation time by on-call engineers.

1 Introduction

Traditional outage detection systems treat multiple signal types in isolation, using separate models for time-series anomalies, alert correlation, support-request spikes etc.. This siloed approach can miss critical cross-signal interactions such as subtle metric drifts that precede alert bursts or support-ticket spikes that occur without corresponding time-series anomalies, leading to delayed or incomplete outage attribution [1]. Jointly modeling service level indicator (SLI) time series, monitor alerts, support-request signals, and historical incident context addresses these limitations by integrating heterogeneous evidence into a single inference path, enabling earlier detection, reducing false positives from single-signal noise, and producing more faithful causal attributions. Our proposed framework reimagines this approach by unifying diverse data modalities as time series embeddings and graph nodes within a single retrieval architecture, allowing numerical similarity and symbolic reasoning to operate together for richer, cross-signal outage detection.

2 Methods

The model operates in a two-components primarily. Firstly converting SLI time series into unified embedded vectors that support similarity search. These signals are transformed into feature embeddings because these representations capture temporal dynamics and metric-shape similarities that enable matching against historical outage patterns. Secondly, categorical signals-such as SLI names, monitors and root-causes are encoded for graph-based reasoning, leveraging GraphRAG [2] to model relationships, dependencies and contextual structure across events. Once the model is trained, live signals are cleaned, embedded, and matched against historical patterns through vector search and graph traversal using the learned time series embeddings and the constructed knowledge graph. This multi-modal integration through GraphRAG enables early detection via cross-modal triangulation, improves robustness by weighting evidence across modalities, and produces explainable multi-vidence reasoning paths that guide engineers towards accurate outage attribution.

3 Early Results

In initial experiments our system achieved 84% accuracy and improved recall by 15% over time-series only baselines. The knowledge graph is a key mechanism that links disparate signal types through learned relationships. When different alert events occur alongside std SLIs degradation, graph traversal identifies their historical co-occurrence across past incidents, including how many escalated to outages and which root causes were previously involved. These correlations provide predictive signal unavailable to unimodal methods. Likewise, support-ticket surges connect to incident nodes quantifying typical lead times. Moreover, 23% of the correctly predicted outages in our tests show sub-threshold time series SLI anomalies which make them detectable only through multi source data fusion.

Our framework naturally extends to new signals without architectural changes. For example, log-based signals can be integrated by embedding log event sequences and linking them to incidents. Furthermore, external dependency outages can also be incorporated via new nodes with temporal alignment logic. This extensibility make the proposed Graph-enhanced retrieval approach a future-proof platform for scalable, multi-modal outage detection as observability signals continue to diversify in cloud services.

References

[1] Peipeng Wang, Xiuguo Zhang, Zhiying Cao, and Zihan Chen, MADMM: microservice system anomaly detection via multi-modal data and multi-feature extraction. *Neural Computing and Applications*, 36:15739–15757, 2024.

[2] Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang, Graph Retrieval-Augmented Generation: A Survey. *arXiv preprint arXiv:2408.08921*, 2024.