A Recovery Module for an Incident Management Assistant with Missing and Imbalanced Data

Alisson N. Bonato[1], **Jade S. Hatanaka Marques**[1,2], **Rajnish Kumar**[1]

[1]Amadeus IT Group – Sophia Antipolis - France

[2]Instituto de Informática – Universidade Federal do Rio Grande do Sul (UFRGS)

{alissonbonatto,jadehatanaka}@gmail.com, rajnish.kumar@amadeus.com

Abstract. *In Information Technology (IT) Service Management, modern approaches leveraging Artificial Intelligence for IT Operations (AIOps) typically rely on abundant, high-quality observability data. However, many real-world environments operate under severe data constraints, such as scarce historical records and incomplete information. This paper introduces AIMA+, an incident management assistant designed specifically for these challenging conditions. Given an incident's textual description as input, AIMA+ outputs a ranked list of predicted recovery actions, retrieves similar past incidents, and generates a synthesized textual summary to guide operators. Our core contribution is a methodology that addresses data incompleteness through a Large Language Model (LLM)-based data augmentation strategy and data scarcity with an interpretable, multi-pronged framework. Our experiments show that augmenting the dataset with expert-guided LLM classifications dramatically improved predictive performance, increasing the macro F1-score from a baseline of 0.2 to 0.6. This work presents a pragmatic blueprint for developing effective AIOps decision support tools in realistic, imperfect industrial settings.*

Keywords: Incident Management, AIOps, Recovery Actions, Imbalanced Data, Large Language Models, Multilabel Classification

1. Introduction

Incident management in dynamic IT environments, such as those supporting automated testing and Continuous Integration/Continuous Deployment (CI/CD) pipelines, presents unique challenges. A high velocity of software changes leads to a frequent stream of incidents, requiring manual diagnosis and recovery by operators. The state-of-the-art in Artificial Intelligence for IT Operations (AIOps) offers powerful solutions but is largely predicated on the availability of rich, historical telemetry.

Prominent research like DeepLog [Du 2017] and LogRobust [Zhang 2019] requires vast historical logs to model system behavior, while advanced Root Cause Analysis (RCA) frameworks like NetSieve [Yuan 2021] depend on complete, end-to-end distributed traces. These data-intensive approaches are unsuitable for our operational context, which lacks long-term observability data. Concurrently, a separate research stream addresses data scarcity using techniques like Information Retrieval [Nguyen 2012] or few-shot learning with Large Language Models (LLMs) [Chen 2023]. While more applicable, these methods typically assume that the limited available data is well-structured and complete.

Our environment suffers from a compounded challenge: our data is not merely scarce but also fundamentally incomplete. Our incident dataset contains fewer than 9,000 relevant records, and over 70% of these lack structured recovery action labels, with solutions often documented in non-standard, free-text descriptions, if at all. This paper introduces AIMA+, a holistic framework designed to operate effectively under these dual constraints.

Our contribution is a pragmatic methodology that first repairs the incomplete dataset using LLMs and then applies a multi-pronged, interpretable synthesis approach to provide decision support. Rather than pursuing full automation, AIMA+ is designed as a human-in-the-loop system to augment operator expertise and reduce Mean Time To Resolution (MTTR). This work is primarily focused on the application of AIOps techniques to solve practical challenges in IT incident management.

2. Methodology

Our methodology is structured to first address the data quality issues and then to build an effective decision support system on the repaired dataset.

2.1. Dataset and Preprocessing

The initial dataset comprises approximately 9,000 incident records from an IT Service Management (ITSM) tool. Each record contains a textual description of a failed software script. A preliminary filtering step was performed to remove incidents related to deprecated software and scripts, further reducing the size of the usable dataset.

2.2. Addressing Data Incompleteness: LLM-based Label Augmentation

overcome the issue of missing structured recovery action labels, we developed a data augmentation pipeline to classify the unstructured, free-text descriptions of solutions into one of 14 predefined labels (e.g., *FALLBACK*, *APPLICATION RESTART*). We systematically evaluated three prompting strategies for this LLM-based classification task:

- **Zero-shot Prompting:** The LLM was provided only with the incident text and the list of possible labels.
- **Few-shot Prompting:** The prompt included a small number of correctly classified examples to guide the model's reasoning.
- **Context-Enriched Few-shot Prompting:** The few-shot prompt was enhanced with domain-specific definitions and criteria for each of the 14 labels, as defined by senior incident managers.

This process allowed us to convert thousands of previously unlabeled records into a structured, usable format for training our machine learning model.

2.3. The AIMA+ Decision Support Framework

+ synthesizes information from four distinct modules to provide a holistic, interpretable recommendation to the operator.

Past Similar Incident Ranker This module uses the BM25 algorithm, a robust and efficient Information Retrieval model, to identify and rank past incidents that are textually similar to the current one. This provides operators with historical precedent, a technique validated in prior work on bug report analysis [Nguyen 2012].

Contextual Cues Module Given the absence of historical logs, this module focuses on just-in-time contextual data highly correlated with script failures. It automatically surfaces relevant recent events, primarily focusing on application deployments and configuration changes, which are known to be a leading cause of production incidents [Wang 2019].

Recovery Action Predictor An XGBoost multilabel classifier was trained on the augmented incident descriptions to predict a probabilistic ranking of the 14 potential recovery actions. The final model was carefully tuned: we employed Sequential Feature Selection (SFS), used weighted training in XGBoost to manage class imbalance, optimized hyperparameters through a Random-then-Grid search, and calibrated probability thresholds against the macro F1-score to ensure balanced performance across all classes.

LLM-Powered Synthesis Module The final module acts as a reasoning engine, integrating the outputs from the previous three components. It follows a Retrieval-Augmented Generation (RAG) pattern [Chen 2023], where the "retrieved" context is the structured output of our ranker, contextual cues, and predictor. This module generates a concise, formatted summary that presents a final recommendation and the supporting evidence, designed for rapid consumption by incident managers.

3. Results and Discussion

3.1. Impact of LLM-based Data Augmentation on Predictive Performance

Our experiments demonstrated the critical role of LLM-based data augmentation. A baseline XGBoost classifier trained on the original, non-augmented data with its few labeled examples performed poorly, achieving a micro-average F1-score of 0.3 and a macro-average F1-score of 0.2. A first augmentation pass, using an unrefined LLM prompt, provided a significant improvement, increasing the scores to a micro-average F1 of 0.54 and a macro-average F1 of 0.28. However, after refining this prompt with expert feedback from Incident Managers, the performance of our final tuned model dramatically improved to a **micro-average F1 of 0.7 and a macro-average F1 of 0.6**. This three-fold improvement in the macro F1-score underscores the necessity of not only using LLMs for augmentation but also grounding them in domain expertise.

3.2. User-Centric Evaluation and System Impact

The primary goal of AIMA+ was to improve operational efficiency and operator trust. We conducted a two-week A/B test with ten incident managers.

- **Performance Impact:** The group using AIMA+ demonstrated a **35% reduction in Mean Time To Resolution (MTTR)** compared to the control group using traditional methods.
- **Trust and Usability:** AIMA+ achieved a **System Usability Scale (SUS) score of 88.5**, well above the industry average, indicating excellent usability.

3.3. Discussion

Beyond these quantitative gains, expert feedback confirmed the importance of interpretability. Operators reported that the model's summaries and the similar-case retrieval feature gave them crucial insight into the rationale behind suggested recovery actions,

which accelerated triage and improved their decision confidence. Therefore, AIMA+ not only improves predictive performance under data scarcity but also enhances the entire recovery process through greater transparency and context.

Our results show that in environments constrained by data scarcity and incompleteness, a human-in-the-loop, interpretable system can be more effective than a "black-box" autonomous solution. The modular, multi-pronged design of AIMA+ provides operators with a comprehensive view, which was crucial for building trust. The modular architecture also provides a clear path for future extensions.

4. Conclusion

This work presented AIMA+, an incident management assistant that provides effective decision support in the face of severe data scarcity and incompleteness. By first repairing our dataset using a novel LLM-based augmentation technique and then synthesizing insights from multiple specialized modules, AIMA+ successfully reduces resolution times and earns operator trust. Our work contributes a practical, validated methodology for building AIOps tools that are resilient to the data challenges commonly found in real-world industrial environments.

Future Work The modular nature of AIMA+ allows for straightforward extension. Our immediate next step is to integrate a conversational interface (chatbot) that will allow operators to query the system's components in natural language. In the long term, as more high-quality data is collected through the system, we will explore transitioning parts of the framework to more automated models.

References

Chen, Z., et al. (2023). Chat-based Incident Triage with Retrieval Augmented Generation. In *ESEC/FSE*.

Du, M., et al. (2017). DeepLog: Anomaly Detection and Diagnosis from System Logs through Deep Learning. In *CCS*.

Nguyen, T. T., et al. (2012). Recommending Similar Bug Reports and Their Fixes. In *ICSE*.

Wang, X., et al. (2019). Learning to Predict and Mitigate Push-Caused Production Failures. In *OSDI*.

Yuan, Z., et al. (2021). NetSieve: A Scalable and Robust Causal Inference Framework for Incident Analysis. In *SOSP*.

Zhang, X., et al. (2019). LogRobust: A Robust Online System Log Anomaly Detection System. In *USENIX Security Symposium*.