RADemics

Recurrent Neural Networks and LSTM for Temporal Anomaly Detection in System Logs

Parvathi R, A. Agalya, Ayesha Taranum

SRM Arts And Science College, SRM Institute of science and technology, Computer Science and Engineering

Recurrent Neural Networks and LSTM for Temporal Anomaly Detection in System Logs

[1]Parvathi R, Assistant Professor, Computer Science with Artificial Intelligence, SRM Arts And Science College, Kattankulathur. parvathikrishna@gmail.com

[2]A. Agalya, Assistant Professor, Computing Technologies, SRM Institute of science and technology, Kattankulathur, Chennai. agalya@smist.edu.in

[3]Ayesha Taranum, Associate Professor & Dy. Head, Computer Science and Engineering, Vidyavardhaka College of Engineering, Mysore, Karnataka. ayeshataranum.cs@vvce.ac.in

Abstract

Anomaly detection in system logs is a critical task for maintaining the integrity, security, and performance of modern computing systems. Traditional methods for detecting anomalies in system logs often struggle with the complexities of sequential and time-dependent data. This chapter explores the application of Long Short-Term Memory (LSTM) networks, a powerful class of Recurrent Neural Networks (RNNs), for detecting temporal anomalies in system logs. LSTMs are capable of learning long-term dependencies within sequential data, making them highly effective at identifying complex, time-dependent patterns that signify anomalous events. However, the inherent black-box nature of LSTM models poses challenges for interpretability and transparency, which are crucial for real-world applications, especially in domains requiring high accountability, such as cybersecurity and system monitoring. To address this challenge, the chapter also examines the integration of LSTM-based anomaly detection with rule-based systems, offering a hybrid approach that combines the predictive power of deep learning with the transparency and explainability of traditional rule-based methods. By leveraging activation mapping, feature attribution, and time-series decomposition, the chapter demonstrates how LSTM models can be made more interpretable, enabling users to trace and validate the model's decisions. The integration of these techniques ensures not only the high accuracy of anomaly detection but also provides actionable insights and clearer explanations for system administrators and security analysts. The chapter discusses the advantages and limitations of LSTM models in anomaly detection, emphasizing their ability to adapt to evolving log data while providing robust, real-time monitoring capabilities. Furthermore, it outlines future research directions for enhancing model transparency and interpretability, focusing on the development of hybrid approaches that balance performance and explainability. This work offers a comprehensive framework for utilizing LSTMs in complex, high-dimensional log data environments, paving the way for more reliable and explainable anomaly detection systems in critical infrastructures.

Keywords: Long Short-Term Memory (LSTM), anomaly detection, system logs, temporal dependencies, interpretability, rule-based systems.

Introduction

Anomaly detection in system logs is an essential task for ensuring the proper functioning, security, and reliability of modern computing systems [1]. System logs, which chronicle a series of time-stamped events, provide insights into the operations and behaviors of a system [2]. These logs are invaluable for identifying irregularities that could indicate system failures, security breaches, or performance degradation [3]. As computing systems grow more complex, the volume of log data has increased significantly, making manual or rule-based approaches insufficient for detecting anomalies in real-time. Anomalies in system logs are often subtle and evolve over time, requiring models that can effectively capture long-term dependencies and sequential patterns [4]. Traditional methods, such as statistical techniques and rule-based systems, struggle with these complex, time-dependent structures, and they fail to adapt quickly to the dynamic nature of modern systems [5].

In recent years, Long Short-Term Memory (LSTM) networks, a specialized form of Recurrent Neural Networks (RNNs), have shown considerable promise for sequential anomaly detection, particularly in time-series data [6]. LSTMs are designed to learn temporal dependencies by maintaining an internal memory of past events, which allows them to recognize patterns over long time spans [7]. This is crucial for detecting anomalies that may not be immediately obvious in short-term data but develop over extended periods, such as performance degradation or gradual security breaches [8]. Their ability to retain important information across many time steps enables LSTMs to model the complex relationships within system logs, which often consist of sequences of events that build upon each other [9]. Thus, LSTMs offer a powerful tool for identifying complex, time-dependent anomalies that would be difficult for traditional anomaly detection methods to capture [10].

LSTM models are often considered "black boxes" due to their lack of interpretability [11]. While they excel at pattern recognition, LSTMs do not inherently provide explanations for their decisions [12]. In mission-critical systems such as cybersecurity or system performance monitoring, understanding why an anomaly is flagged is essential for decision-making and trust [13]. Analysts and operators need to know not only that an anomaly has occurred but also the rationale behind its detection [14]. This lack of transparency can hinder the adoption of LSTM-based models in practical applications, where trust and explainability are vital for regulatory compliance and operational accountability. Thus, improving the interpretability of LSTM-based anomaly detection is a significant challenge in deploying these models effectively [15].