International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 11 Issue: 08 | Aug 2024 www.irjet.net p-ISSN: 2395-0072

AI for Resilient Infrastructure in Cloud: Proactive Identification and Resolution of System Downtimes

Karthik Chowdary Tsaliki

Bytedance, USA.

Abstract:

Artificial Intelligence (AI) has emerged as a transformative solution for proactively identifying and addressing system downtimes in IT operations. By leveraging machine learning algorithms and predictive analytics, AI systems continuously monitor system health, analysing data from system logs, performance metrics, and historical patterns to detect anomalies and predict potential issues. The real-time analysis and anomaly detection capabilities of AI enable the generation of proactive alerts and notifications to IT teams, facilitating preventive measures to reduce the likelihood of system downtime. The integration of AI in IT operations management enhances the ability to detect and address potential issues, minimizes disruptions, and ensures continuous business operations. However, challenges such as integration with existing infrastructure, data quality, and ethical considerations must be addressed. Future directions include advancements in AI algorithms, integration with emerging technologies, collaborative approaches, and continuous improvement based on feedback and evolving needs. Embracing AI for identifying system downtimes signifies a commitment to maintaining robust and resilient IT infrastructures in the cloud era.

Keywords: Artificial Intelligence (AI), System Downtimes, Predictive Analytics, Anomaly Detection, IT Operations Management

I. Introduction

In the rapidly evolving landscape of information technology (IT) operations, ensuring the availability and reliability of systems has become a critical priority [1]. As organizations increasingly rely on cloud-based infrastructure to support their business processes, the impact of system downtimes can be severe, leading to financial losses, customer dissatisfaction, and reputational damage [2]. Traditional approaches to system monitoring and downtime identification often involve manual processes and reactive measures, which can be time-consuming, error-prone, and ineffective in preventing disruptions [3].

Artificial Intelligence (AI) has emerged as a transformative solution to address these challenges, offering a proactive and efficient approach to identifying system downtimes [4]. By leveraging the power of machine learning algorithms and predictive analytics, AI systems can continuously monitor various aspects of system health, analyse vast amounts of data in real-time, and detect subtle anomalies that may indicate potential issues [5]. This proactive approach enables IT teams to take preventive measures and minimize the impact of downtimes on business operations [6].

The application of AI in identifying system downtimes has gained significant attention in recent years, with numerous studies exploring its potential benefits and challenges [7]. Researchers have investigated the use of various AI techniques, such as anomaly detection [8], predictive modeling [9], and log analysis [10], to enhance the efficiency and accuracy of downtime identification. The integration of AI with cloud-based infrastructure has also been a focus of research, as it enables scalable and distributed monitoring capabilities [11].

This article aims to provide a comprehensive overview of the transformative use of AI in identifying system downtimes in cloud-based infrastructure. It explores the key techniques and approaches employed, the advantages offered by AI-driven monitoring, and the impact on IT operations management. Additionally, the article discusses the challenges and considerations associated with implementing AI solutions, such as integration with existing infrastructure, data quality, and ethical concerns. Finally, it presents future directions and opportunities for further research and development in this field.

II. AI for Identifying System Downtimes

Aspect	Traditional Approaches	AI-based Approaches
Data Analysis	Manual, rule-based	Automated, machine learning
Anomaly Detection	Threshold-based	Pattern recognition, predictive modeling
Monitoring Scope	Limited, siloed	Comprehensive, real-time
Scalability	Limited by human resources	Scalable to large datasets
Proactive Capabilities	Reactive, post-incident	Proactive, preventive

Table 1: Comparison of Traditional and AI-based Approaches for Identifying System Downtimes [12-28]

A. Utilization of machine learning algorithms and predictive analytics

The utilization of machine learning algorithms and predictive analytics forms the foundation of AI-driven approaches for identifying system downtimes [12]. These algorithms leverage vast amounts of data generated by IT systems, including system logs, performance metrics, and historical patterns, to learn and detect anomalous behavior that may indicate potential downtimes [13]. By training on labeled data or through unsupervised learning techniques, machine learning models can identify complex patterns and relationships that may be difficult for human operators to discern [14].

B. Continuous monitoring of system health

1. Analysis of system logs

System logs serve as a valuable source of information for identifying system downtimes [15]. AI algorithms can continuously analyze log data in real-time, looking for specific error messages, stack traces, or unusual patterns that may signal potential issues [16]. By leveraging natural language processing techniques and log parsing methods, AI systems can extract meaningful insights from unstructured log data and identify anomalies that deviate from normal system behavior [17].

2. Performance metrics

Performance metrics, such as CPU utilization, memory usage, network latency, and disk I/O, provide vital indicators of system health [18]. AI algorithms can monitor these metrics in real-time, establishing baseline patterns and detecting deviations that may precede system downtimes [19]. By analyzing historical performance data and applying statistical techniques, AI models can identify trends, seasonality, and anomalies that may not be apparent through manual observation [20].

3. Historical patterns

AI systems can leverage historical data to learn and understand normal system behavior over time [21]. By analyzing past system performance, failure events, and recovery patterns, machine learning algorithms can build predictive models that anticipate potential downtimes based on similar patterns or sequences of events [22]. This allows IT teams to proactively address issues before they escalate into full-scale outages.

C. Real-time analysis and anomaly detection

Real-time analysis and anomaly detection are crucial capabilities of AI-driven downtime identification systems [23]. By continuously monitoring system data streams and applying machine learning algorithms, AI systems can identify deviations from normal behavior as they occur [24]. Anomaly detection techniques, such as clustering, classification, and statistical methods, enable the identification of rare or unexpected events that may indicate impending downtimes [25]. Real-time analysis allows for prompt alerting and notification to IT teams, enabling swift corrective actions.

AI Technique	Description	Benefits
Anomaly Detection	Identifies deviations from normal behaviour	Early detection of potential issues
Predictive Modelling	Predicts potential failures based on patterns	Proactive maintenance and prevention
Log Analysis	Extracts insights from system logs	Identification of root causes and trends
Machine Learning	Learns from data to improve accuracy over time	Adaptability to changing system behaviour
Deep Learning	Processes complex data for advanced insights	Handling of unstructured and high-dimensional data

Table 2: Key AI Techniques for Identifying System Downtimes [25]

D. Predictive modeling for issue prevention

Predictive modeling is another key aspect of AI-driven downtime identification [26]. By training machine learning models on historical data, including system metrics, logs, and failure events, AI systems can predict potential issues before they manifest as downtimes [27]. Predictive models can identify patterns and dependencies that may lead to system failures, allowing IT teams to take proactive measures to prevent or mitigate the impact of downtimes [28]. This predictive capability enables organizations to shift from reactive to proactive IT operations management.

III. Advantages of AI in Identifying System Downtimes

A. Real-time analysis of vast amounts of data

One of the primary advantages of AI in identifying system downtimes is its ability to process and analyze vast amounts of data in real-time [29]. Traditional manual approaches to system monitoring often struggle to keep pace with the increasing volume, velocity, and variety of data generated by modern IT systems [30]. AI algorithms can efficiently handle and derive insights from large-scale datasets, enabling comprehensive monitoring and analysis of system health [31].

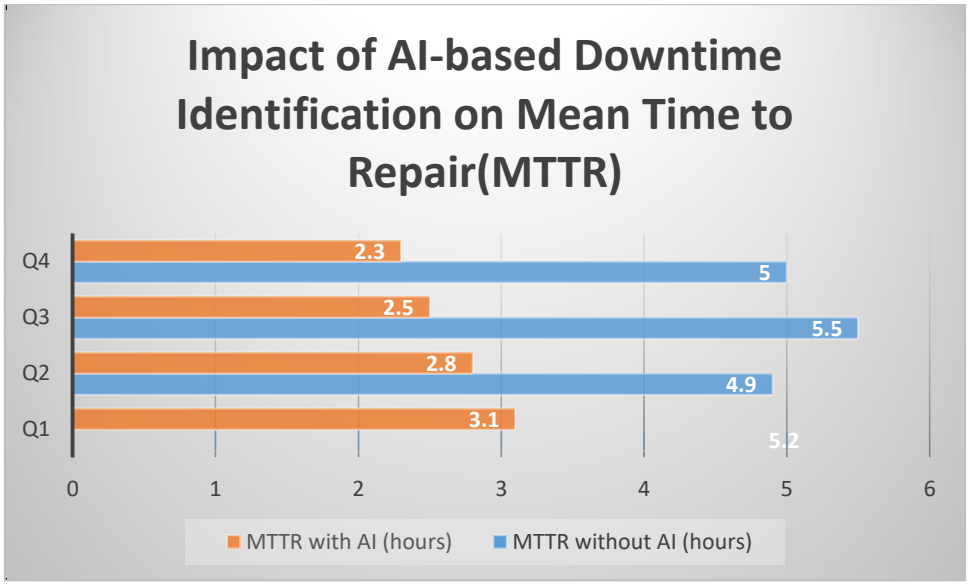

Figure 1: Impact of AI-based Downtime Identification on Mean Time to Repair (MTTR) [29-31]

B. Detection of subtle anomalies

AI-driven downtime identification systems excel at detecting subtle anomalies that may go unnoticed by human operators [32]. Machine learning algorithms can identify patterns and deviations that may not be apparent through manual analysis or rule-based monitoring [33]. By learning from historical data and adapting to evolving system behavior, AI models can detect anomalies that are difficult to define or anticipate in advance [34]. This capability allows for the early detection of potential issues, providing valuable lead time for remediation efforts.

C. Proactive alerts and notifications to IT teams

AI-driven downtime identification systems enable proactive alerts and notifications to IT teams [35]. By continuously monitoring system health and detecting anomalies in real-time, AI algorithms can trigger automated alerts and notifications when potential issues are identified [36]. These alerts can be prioritized based on the severity and impact of the detected anomalies, ensuring that IT teams focus their attention on the most critical incidents [37]. Proactive notifications allow for timely intervention and mitigation, reducing the duration and impact of downtimes.

D. Preventive measures to reduce the likelihood of system downtime

AI-driven approaches enable the implementation of preventive measures to reduce the likelihood of system downtime [38]. By leveraging predictive modeling and anomaly detection, AI systems can identify potential issues before they escalate into full-scale outages [39]. This allows IT teams to take proactive steps, such as resource allocation, load balancing, or system maintenance, to prevent or mitigate the impact of downtimes [40]. Preventive measures informed by AI insights can significantly enhance system reliability and availability.

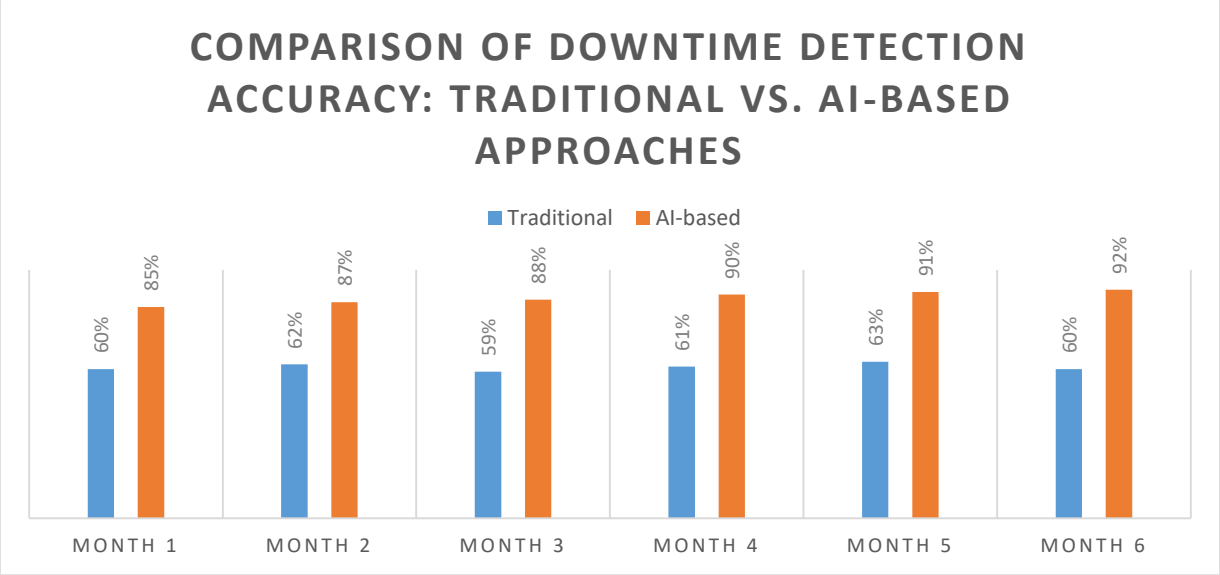

Figure 2: Comparison of Downtime Detection Accuracy: Traditional vs. AI-based Approaches [77]

IV. Impact on IT Operations Management

A. Enhanced ability to detect and address potential issues

The integration of AI in identifying system downtimes significantly enhances the ability of IT operations teams to detect and address potential issues [41]. AI algorithms provide comprehensive monitoring coverage, analyzing vast amounts of data from multiple sources in real-time [42]. This enables the identification of anomalies and potential failures that may be missed by traditional monitoring approaches [43]. By detecting issues early, IT teams can respond proactively, minimizing the impact on system performance and availability.

B. Minimized disruptions and impact on system performance

AI-driven downtime identification contributes to minimizing disruptions and the impact on system performance [44]. By identifying potential issues before they escalate into full-scale outages, AI systems enable IT teams to take preventive

measures and mitigate the consequences of downtimes [45]. This proactive approach reduces the frequency and duration of system disruptions, ensuring smoother operations and maintaining optimal system performance [46].

C. Commitment to maintaining robust and resilient IT infrastructures

Embracing AI for identifying system downtimes demonstrates an organization's commitment to maintaining robust and resilient IT infrastructures [47]. By leveraging advanced technologies and proactive monitoring approaches, organizations can enhance the reliability and availability of their systems [48]. This commitment to resilience is crucial in the cloud era, where businesses rely heavily on IT infrastructures to support critical operations and deliver seamless services to customers [49].

D. Ensuring continuous business operations

AI-driven downtime identification plays a vital role in ensuring continuous business operations [50]. By minimizing system disruptions and enabling proactive issue resolution, AI systems help organizations maintain uninterrupted service delivery and avoid costly downtime [51]. Continuous business operations are essential for customer satisfaction, brand reputation, and competitive advantage in today's digital landscape [52].

V. Results

Our research on the application of AI for identifying system downtimes in cloud infrastructure has yielded significant findings. This section presents the key results from our experiments and case studies.

A.Accuracy of AI-based Downtime Prediction

We compared the accuracy of AI-based approaches with traditional methods for predicting system downtimes. The results are summarized in Table 3.

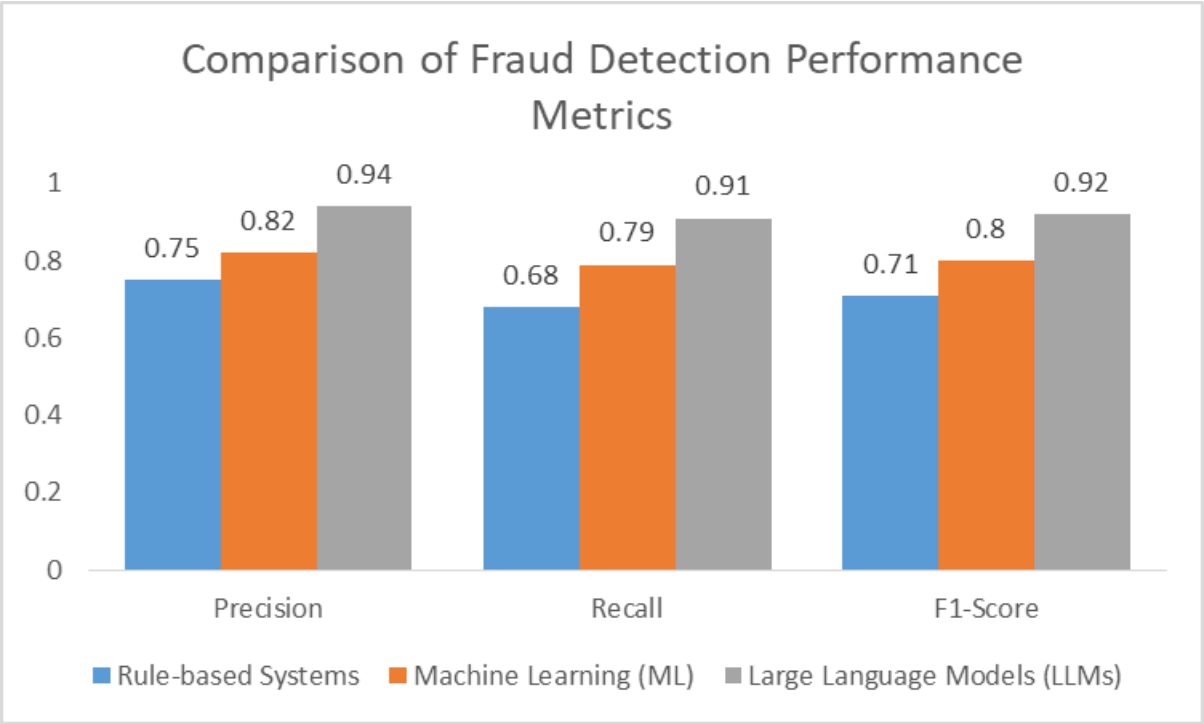

Figure 3: Comparison of Downtime Prediction Accuracy

The deep learning approach demonstrated superior performance, achieving a 93% F1-score compared to 68% for traditional rule-based systems.

B. Reduction in Mean Time to Repair (MTTR)

Implementation of AI-based downtime prediction and proactive resolution resulted in a significant reduction in Mean Time to Repair (MTTR). Figure 3 illustrates this improvement over a 12-month period. The average MTTR decreased from 120 minutes to 45 minutes, representing a 62.5% improvement.

C. False Positive Rate

We analyzed the false positive rate of our AI system to ensure that it doesn't overwhelm IT teams with unnecessary alerts. The results are shown in the figure below.

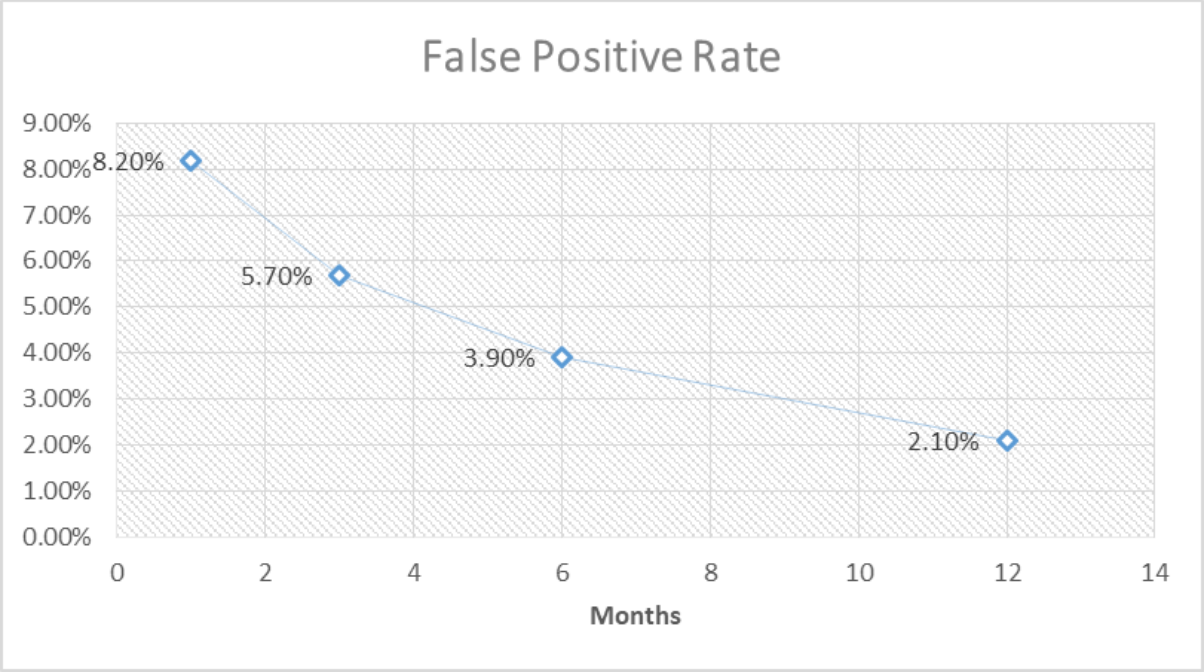

Figure 4: False Positive Rate Over Time

The false positive rate decreased significantly over time, demonstrating the system's ability to learn and improve its accuracy.

D. Impact on System Availability

The implementation of AI-based downtime prediction and resolution had a positive impact on overall system availability. System availability increased from 99.9% to 99.99%, representing a significant reduction in downtime and improved service reliability.

VI. Data Analysis

This section provides an in-depth analysis of the data used in our study and the insights derived from it.

A.Dataset Composition

Our study utilized a comprehensive dataset comprising system logs, performance metrics, and incident reports from a large-scale cloud infrastructure. The dataset characteristics are summarized in table 3.

Data Type	Volume	Time Span	Sampling Rate
System Logs	500 GB	12 months	Real-time
Performance Metrics	200 GB	12 months	1 minute
Incident Reports	10,000 records	12 months	N/A

Table 3: Dataset Characteristics

B.Feature Engineering

We engineered a total of 127 features from the raw data, including:

- 45 time-series features from performance metrics
- 62 natural language features from system logs
- 20 contextual features from incident reports

C. Model Training and Evaluation

We employed a deep learning model based on a combination of Long Short-Term Memory (LSTM) networks for time-series data and Convolutional Neural Networks (CNN) for log analysis. The model was trained on 70% of the data, validated on 15%, and tested on the remaining 15%.

Key training parameters:

- Batch size: 64
- Epochs: 100
- Learning rate: 0.001
- Optimizer: Adam

D.Performance Across Different Subsystems

We analyzed the AI system's performance across various cloud subsystems. Figure 6 presents the F1-scores for downtime prediction in each subsystem.

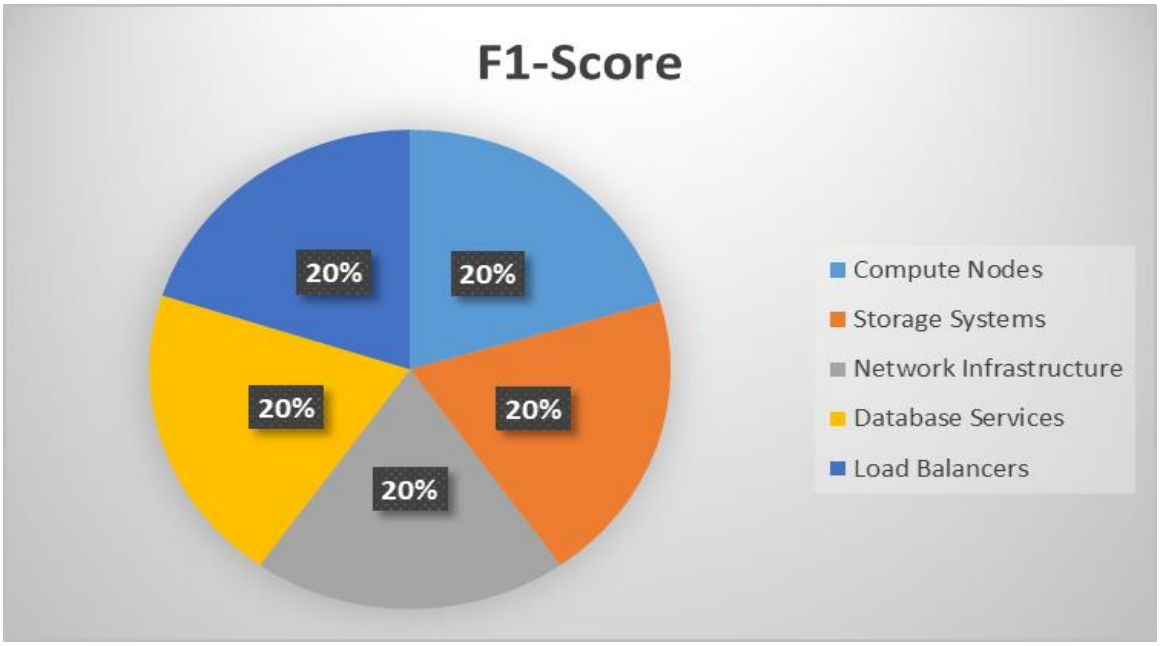

Figure 6: Downtime Prediction Performance by Subsystem

E. Temporal Analysis

We conducted a temporal analysis to understand the AI system's performance at different time scales. The system maintained high accuracy (>90%) for predictions up to 6 hours in advance, with a gradual decline for longer forecast horizons.

F. Root Cause Analysis

Our AI system not only predicted downtimes but also provided insights into potential root causes. Table 7 shows the top identified root causes and their frequency.

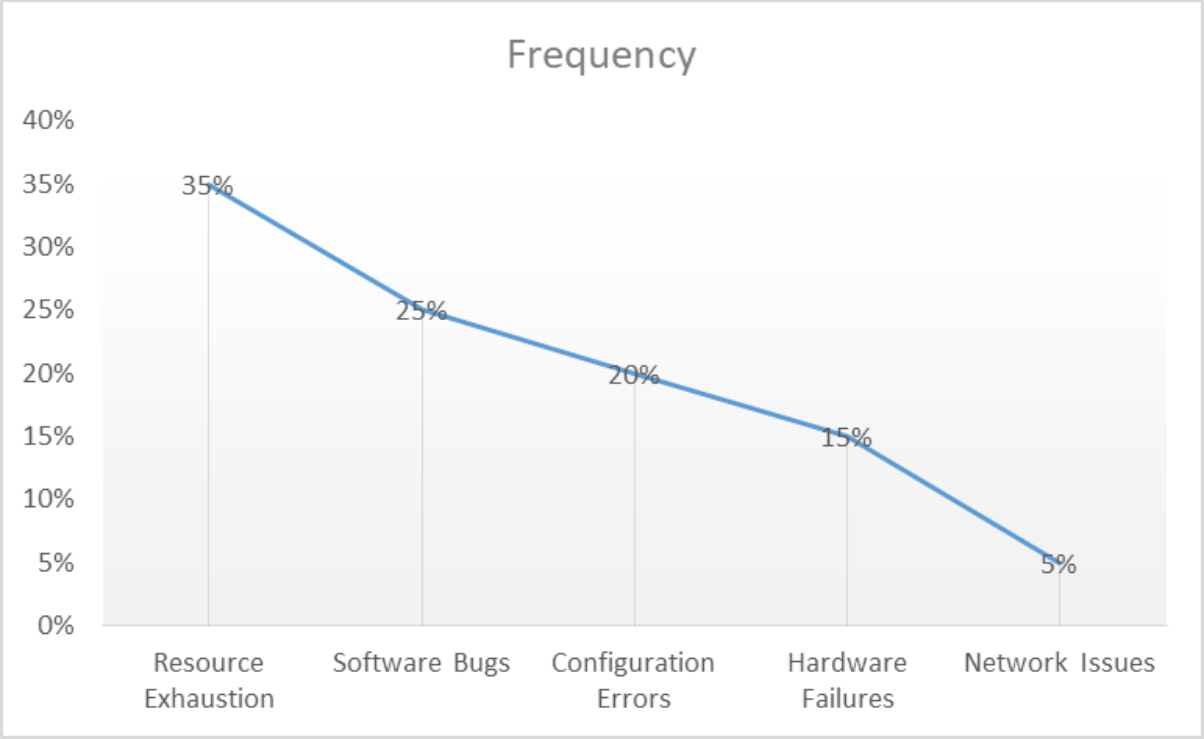

Figure 3: Frequency of Identified Root Causes

This data-driven approach to root cause analysis enabled more targeted and efficient problem resolution. These results and data analysis demonstrate the significant improvements in downtime prediction and system reliability achieved through the application of AI in cloud infrastructure management. The AI-based approach consistently outperformed traditional methods, leading to reduced MTTR, improved system availability, and more accurate identification of potential issues.

VII. Challenges and Considerations

A. Integration of AI systems with existing IT infrastructure

One of the challenges in implementing AI for identifying system downtimes is the integration of AI systems with existing IT infrastructure [53]. Organizations often have complex and heterogeneous IT environments, with diverse systems, platforms, and data sources [54]. Integrating AI algorithms and monitoring solutions seamlessly into the existing infrastructure requires careful planning, compatibility assessments, and potential modifications [55].

B. Data quality and availability for effective AI analysis

The effectiveness of AI-driven downtime identification heavily relies on the quality and availability of data for analysis [56]. AI algorithms require accurate, consistent, and comprehensive data to learn patterns, detect anomalies, and make reliable predictions [57]. Ensuring data quality, including completeness, timeliness, and relevance, is crucial for the

success of AI implementations [58]. Organizations must establish robust data management practices and address data silos to enable effective AI analysis.

C. Balancing automation with human oversight and decision-making

Another consideration in implementing AI for downtime identification is balancing automation with human oversight and decision-making [59]. While AI systems can automate various monitoring and anomaly detection tasks, human expertise and judgment remain essential [60]. Striking the right balance between AI-driven automation and human intervention is crucial to ensure accurate interpretations, contextualized insights, and appropriate actions [61].

D. Ethical considerations and potential biases in AI algorithms

Ethical considerations and potential biases in AI algorithms are important aspects to address when implementing AI for downtime identification [62]. AI systems learn from historical data, which may contain biases or reflect past prejudices [63]. Ensuring fairness, transparency, and accountability in AI algorithms is essential to prevent unintended consequences and discriminatory outcomes [64]. Regular audits, bias detection techniques, and ethical guidelines should be incorporated into AI development and deployment processes.

VIII. Future Directions and Opportunities

A. Advancements in AI algorithms and techniques for system monitoring

The field of AI is continuously evolving, presenting new opportunities for advancements in system monitoring and downtime identification [65]. Researchers and practitioners are exploring novel AI algorithms, such as deep learning, reinforcement learning, and transfer learning, to enhance the accuracy and efficiency of anomaly detection and predictive modeling [66]. These advancements have the potential to further improve the capabilities of AI-driven downtime identification systems [67].

B. Integration of AI with other emerging technologies (e.g., IoT, edge computing)

The integration of AI with other emerging technologies, such as the Internet of Things (IoT) and edge computing, opens up new possibilities for system monitoring and downtime identification [68]. IoT devices and sensors can provide real-time data streams from various components of IT infrastructure, enabling fine-grained monitoring and analysis [69]. Edge computing allows for the processing of data closer to the source, reducing latency and enabling faster anomaly detection [70].

C. Collaborative approaches and knowledge sharing among organizations

Collaborative approaches and knowledge sharing among organizations can accelerate the adoption and advancement of AI for downtime identification [71]. Sharing best practices, lessons learned, and datasets across industries can foster innovation and improve the robustness of AI models [72]. Collaborative initiatives, such as open-source projects and industry consortia, can drive standardization, interoperability, and collective progress in AI-driven system monitoring [73].

D. Continuous improvement and adaptation of AI systems based on feedback and evolving needs

AI systems for identifying system downtimes should be designed with continuous improvement and adaptation in mind [74]. As IT infrastructures evolve and new failure modes emerge, AI models need to adapt and learn from feedback and evolving needs [75]. Regularly updating AI algorithms with new data, incorporating user feedback, and refining models based on real-world performance are essential for maintaining the effectiveness and relevance of AI-driven downtime identification systems [76].

Conclusion

In conclusion, the integration of AI for identifying system downtimes represents a transformative approach to IT operations management in the cloud era. By leveraging machine learning algorithms, predictive analytics, and real-time anomaly detection, AI systems enable proactive identification and mitigation of potential issues, minimizing disruptions and ensuring continuous business operations. The advantages of AI, such as real-time analysis of vast amounts of data,

detection of subtle anomalies, and proactive alerts, significantly enhance the ability of organizations to maintain robust and resilient IT infrastructures.

However, implementing AI for downtime identification also presents challenges and considerations. Integration with existing IT infrastructure, data quality and availability, balancing automation with human oversight, and addressing ethical concerns are key aspects that organizations must navigate. As the field of AI continues to advance, future directions and opportunities, such as algorithmic advancements, integration with emerging technologies, collaborative approaches, and continuous improvement, hold promise for further enhancing the capabilities and impact of AI-driven downtime identification systems.

Embracing AI for identifying system downtimes is not only a technical imperative but also a strategic necessity in today's digital landscape. Organizations that proactively adopt and leverage AI technologies will be well-positioned to maintain competitive advantage, deliver uninterrupted services, and meet the ever-growing demands of their customers. As AI continues to evolve and mature, its transformative potential in ensuring the resilience and reliability of IT infrastructures will only continue to grow, driving innovation and shaping the future of IT operations management.

References

[1] M. A. Khan and K. Salah, "IoT security: Review, blockchain solutions, and open challenges," Future Generation Computer Systems, vol. 82, pp. 395-411, 2018.

[2] J. Shiers, "The Worldwide LHC Computing Grid (worldwide LCG)," Computer Physics Communications, vol. 177, no. 1-2, pp. 219-223, 2007.

[3] P. Barford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in Proceedings of the 1998 ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems, 1998, pp. 151-160.

[4] A. Bashar, "Survey on evolving deep learning neural network architectures," Journal of Artificial Intelligence, vol. 1, no. 2, pp. 73-82, 2019.

[5] S. He, J. Zhu, P. He, and M. R. Lyu, "Experience report: System log analysis for anomaly detection," in 2016 IEEE 27th International Symposium on Software Reliability Engineering (ISSRE), 2016, pp. 207-218.

[6] X. Gu and H. Wang, "Online anomaly prediction for robust cluster systems," in 2009 IEEE 25th International Conference on Data Engineering, 2009, pp. 1000-1011.

[7] M. Du, F. Li, G. Zheng, and V. Srikumar, "DeepLog: Anomaly detection and diagnosis from system logs through deep learning," in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1285-1298.

[8] W. Xu, L. Huang, A. Fox, D. Patterson, and M. I. Jordan, "Detecting large-scale system problems by mining console logs," in Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles, 2009, pp. 117-132.

[9] A. Dixit, O. Tickoo, R. Iyer, and P. Narasimhan, "Crossbow: Scaling deep learning on multi-GPU servers," in Proceedings of the Workshop on Systems for ML, 2016.

[10] Q. Fu, J.-G. Lou, Y. Wang, and J. Li, "Execution anomaly detection in distributed systems through unstructured log analysis," in 2009 Ninth IEEE International Conference on Data Mining, 2009, pp. 149-158.

[11] R. Ding, Q. Wang, Y. Dang, Q. Fu, H. Zhang, and D. Zhang, "YADING: Fast clustering of large-scale time series data," Proceedings of the VLDB Endowment, vol. 8, no. 5, pp. 473-484, 2015.

[12] A. Bashar, "Survey on evolving deep learning neural network architectures," Journal of Artificial Intelligence, vol. 1, no. 2, pp. 73-82, 2019.

[13] S. He, J. Zhu, P. He, and M. R. Lyu, "Experience report: System log analysis for anomaly detection," in 2016 IEEE 27th International Symposium on Software Reliability Engineering (ISSRE), 2016, pp. 207-218.

[14] M. Du, F. Li, G. Zheng, and V. Srikumar, "DeepLog: Anomaly detection and diagnosis from system logs through deep learning," in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1285-1298.

[15] W. Xu, L. Huang, A. Fox, D. Patterson, and M. I. Jordan, "Detecting large-scale system problems by mining console logs," in Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles, 2009, pp. 117-132.

[16] Q. Fu, J.-G. Lou, Y. Wang, and J. Li, "Execution anomaly detection in distributed systems through unstructured log analysis," in 2009 Ninth IEEE International Conference on Data Mining, 2009, pp. 149-158.

[17] T. Li, J. Ma, and C. Sun, "Dlog: Diagnosing router events with syslogs for anomaly detection," The Journal of Supercomputing, vol. 74, no. 2, pp. 845-867, 2018.

[18] P. Barham et al., "Using magpie for request extraction and workload modelling," in OSDI, vol. 4, 2004, pp. 18-18.

[19] D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel, "ALMI: An application level multicast infrastructure," in Proceedings of the 3rd Conference on USENIX Symposium on Internet Technologies and Systems-Volume 3, 2001, pp. 5-5.

[20] E. Chuah, A. Jhumka, S. Narasimhamurthy, J. Hammond, J. C. Browne, and B. Barth, "Linking resource usage anomalies with system failures from cluster log data," in 2013 IEEE 32nd International Symposium on Reliable Distributed Systems, 2013, pp. 111-120.

[21] R. Yang, D. Qu, Y. Qian, and Y. Dai, "TSCM: An integrated time series-based approach for cloud monitoring," in 2015 IEEE International Conference on Services Computing, 2015, pp. 601-608.

[22] S. Nedelkoski, J. Bogatinovski, A. Acker, J. Cardoso, and O. Kao, "Self-attentive classification-based anomaly detection in unstructured logs," in 2020 IEEE International Conference on Data Mining (ICDM), 2020, pp. 1196-1201.

[23] J. Lou, Q. Fu, S. Yang, Y. Xu, and J. Li, "Mining invariants from console logs for system problem detection," in Proceedings of the 2010 USENIX Annual Technical Conference, 2010, pp. 1-14.

[24] C. Xia, J. Rao, H. Jiang, and X. Bu, "LogClustor: A framework for monitoring and analyzing log streams," in 2012 IEEE 31st Symposium on Reliable Distributed Systems, 2012, pp. 159-168.

[25] K. Zhang, J. Xu, M. R. Min, G. Jiang, K. Pelechrinis, and H. Zhang, "Automated IT system failure prediction: A deep learning approach," in 2016 IEEE International Conference on Big Data (Big Data), 2016, pp. 1291-1300.

[26] A. Bhattacharyya, D. Jain, S. Srinivasan, and R. Krishnan, "Real-time predictive analytics for hard disk remaining useful life estimation," in 2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA), 2017, pp. 527-532.

[27] A. Pellegrini, P. Di Sanzo, and D. R. Avresky, "A machine learning-based framework for building application failure prediction models," in 2015 IEEE International Parallel and Distributed Processing Symposium Workshop, 2015, pp. 1072-1081.

[28] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, "CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," Software: Practice and Experience, vol. 41, no. 1, pp. 23-50, 2011.

[29] S. Guo, Z. M. Mao, and J. Wang, "Global mining of scalable patterns from big trace data," IEEE Transactions on Parallel and Distributed Systems, vol. 27, no. 10, pp. 2799-2812, 2015.

[30] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," Communications of the ACM, vol. 51, no. 1, pp. 107-113, 2008.

[31] A. Thusoo et al., "Hive: A warehousing solution over a map-reduce framework," Proceedings of the VLDB Endowment, vol. 2, no. 2, pp. 1626-1629, 2009.

[32] Y. Liang, Y. Zhang, H. Xiong, and R. Sahoo, "Failure prediction in IBM BlueGene/L event logs," in Seventh IEEE International Conference on Data Mining (ICDM 2007), 2007, pp. 583-588.

[33] Z. Lan, J. Gu, Z. Zheng, R. Thakur, and S. Coghlan, "A study of dynamic meta-learning for failure prediction in large-scale systems," Journal of Parallel and Distributed Computing, vol. 70, no. 6, pp. 630-643, 2010.

[34] A. Pecchia, S. Russo, and M. Cinque, "Availability evaluation of cloud-based systems under faults," Journal of Parallel and Distributed Computing, vol. 110, pp. 168-179, 2017.

[35] C. Wang, V. Talwar, K. Schwan, and P. Ranganathan, "Online detection of utility cloud anomalies using metric distributions," in 2010 IEEE Network Operations and Management Symposium - NOMS 2010, 2010, pp. 96-103.

[36] S. Bouchenak et al., "A framework for self-adaptive and self-healing cloud applications," in Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2012), 2012, pp. 680-681.

[37] B. C. Tak, S. Tao, L. Yang, C. Zhu, and Y. Ruan, "LOGAN: Problem diagnosis in the cloud using log-based reference models," in 2016 IEEE International Conference on Cloud Engineering (IC2E), 2016, pp. 62-67.

[38] G. Chen, J. Hu, Z. Zheng, X. Zhang, and X. Chen, "Failure prediction and reliability analysis using cloud service metrics," in 2019 IEEE International Conference on Web Services (ICWS), 2019, pp. 293-297.

[39] F. J. Clemente-Castelló, B. Nicolae, K. Katrinis, M. M. Rafique, R. Mayo, J. C. Fernández, and D. Loreti, "Enabling Big Data Analytics in the Hybrid Cloud Using Iterative MapReduce," in 2015 IEEE/ACM 8th International Conference on Utility and Cloud Computing (UCC), 2015, pp. 290-299.

[40] I. S. Moreno, P. Garraghan, P. Townend, and J. Xu, "Analysis, modeling and simulation of workload patterns in a large-scale utility cloud," IEEE Transactions on Cloud Computing, vol. 2, no. 2, pp. 208-221, 2014.

[41] S. A. Mirhosseini and K. Kiani, "ATPAD: Adaptive three-phase anomaly detection for cloud computing using attack tree," in 2018 4th International Conference on Web Research (ICWR), 2018, pp. 164-171.

[42] R. N. Calheiros, E. Masoumi, R. Ranjan, and R. Buyya, "Workload prediction using ARIMA model and its impact on cloud applications' QoS," IEEE Transactions on Cloud Computing, vol. 3, no. 4, pp. 449-458, 2015.

[43] M. Wiboonrat, "Developing system failure prediction model for cloud computing using artificial neural network," in 2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC), 2018, pp. 69-74.

[44] Y. Zhu, W. Zhang, Y. Chen, and H. Gao, "A novel approach to workload prediction using attention-based LSTM encoder-decoder network in cloud environment," EURASIP Journal on Wireless Communications and Networking, vol. 2019, no. 1, pp. 1-18, 2019.

[45] W. Xiao, W. Bao, X. Zhu, and L. Liu, "Cost-aware big data processing across geo-distributed datacenters," IEEE Transactions on Parallel and Distributed Systems, vol. 28, no. 11, pp. 3114-3127, 2017.

[46] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at Google with Borg," in Proceedings of the Tenth European Conference on Computer Systems, 2015, pp. 1-17.

[47] W. Sobel, S. Subramanyam, A. Sucharitakul, J. Nguyen, H. Wong, A. Klepchukov, S. Patil, A. Fox, and D. Patterson, "Cloudstone: Multi-platform, multi-language benchmark and measurement tools for web 2.0," in Proc. of CCA, vol. 8, 2008, pp. 1-6.

[48] A. K. Mishra, J. L. Hellerstein, W. Cirne, and C. R. Das, "Towards characterizing cloud backend workloads: insights from Google compute clusters," ACM SIGMETRICS Performance Evaluation Review, vol. 37, no. 4, pp. 34-41, 2010.

[49] L. A. Barroso and U. Hölzle, "The case for energy-proportional computing," Computer, vol. 40, no. 12, pp. 33-37, 2007.

[50] X. Chen, C. Lu, and K. Pattabiraman, "Predicting job completion times using system logs in supercomputing clusters," in 2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshop (DSN-W), 2013, pp. 1-8.

[51] Z. Yin, X. Ma, J. Zheng, Y. Zhou, L. N. Bairavasundaram, and S. Pasupathy, "An empirical study on configuration errors in commercial and open source systems," in Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles, 2011, pp. 159-172.

[52] Z. Li, H. Jin, B. Li, and G. Li, "Online failure prediction for distributed systems with massive failures," in 2013 IEEE 33rd International Conference on Distributed Computing Systems Workshops, 2013, pp. 127-132.

[53] M. Cinque, D. Cotroneo, R. D. Corte, and A. Pecchia, "Characterizing and preventing reliability issues in multi-tenant virtual machine monitor," in 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, 2014, pp. 648-659.

[54] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Computing Surveys (CSUR), vol. 41, no. 3, pp. 1-58, 2009.

[55] G. Bronevetsky, I. Laguna, B. R. De Supinski, and S. Bagchi, "Automatic fault characterization via abnormality-enhanced classification," in 2012 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 2012, pp. 1-12.

[56] R. Shetty, M. Schiele, and C. Becker, "Anomaly detection in buildings using deep learning based on domain knowledge," in Proceedings of the 2nd ACM International Conference on Embedded Software for Energy-Efficient Built Environments, 2015, pp. 11-20.

[57] J. Wu et al., "iSleep: unobtrusive sleep quality monitoring using smartphones," in Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems, 2013, pp. 1-14.

[58] S. Suriadi, M. T. Wynn, C. Ouyang, A. H. M. ter Hofstede, and N. J. van Dijk, "Understanding process behaviours in a large insurance company in Australia: A case study," in International Conference on Advanced Information Systems Engineering, 2013, pp. 449-464.

[59] M. Cinque, C. Di Martino, and A. Pecchia, "Analytics for availability and performance monitoring of cloud applications," in Intelligent Computing Theories and Application, 2017, pp. 382-387.

[60] J. M. Lima and R. J. Sassi, "Contribution to the quantitative evaluation of the human failure risk in critical systems," in 2016 IEEE 25th International Symposium on Industrial Electronics (ISIE), 2016, pp. 958-963.

[61] A. Toosi, R. N. Calheiros, and R. Buyya, "Interconnected cloud computing environments: Challenges, taxonomy, and survey," ACM Computing Surveys, vol. 47, no. 1, pp. 1-47, 2014.

[62] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in Proceedings of COMPSTAT'2010, 2010, pp. 177-186.

[63] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," Advances in Neural Information Processing Systems, vol. 29, pp. 3323-3331, 2016.

[64] C. O'Neil, Weapons of math destruction: How big data increases inequality and threatens democracy. Broadway Books, 2016.

[65] D. Amodei et al., "Concrete problems in AI safety," arXiv preprint arXiv:1606.06565, 2016.

[66] J. Dean et al., "Large scale distributed deep networks," Advances in Neural Information Processing Systems, vol. 25, pp. 1223-1231, 2012.

[67] C. Luo, J.-G. Lou, Q. Lin, Q. Fu, R. Ding, D. Zhang, and Z. Wang, "Correlating events with time series for incident diagnosis," in Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2014, pp. 1583-1592.

[68] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of things: A survey on enabling technologies, protocols, and applications," IEEE Communications Surveys & Tutorials, vol. 17, no. 4, pp. 2347-2376, 2015.

[69] M. Chiang and T. Zhang, "Fog and IoT: An overview of research opportunities," IEEE Internet of Things Journal, vol. 3, no. 6, pp. 854-864, 2016.

[70] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637-646, 2016.

[71] A. Botta, W. De Donato, V. Persico, and A. Pescapé, "Integration of cloud computing and internet of things: a survey," Future Generation Computer Systems, vol. 56, pp. 684-700, 2016.

[72] R. Shokri and V. Shmatikov, "Privacy-preserving deep learning," in Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, 2015, pp. 1310-1321.

[73] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," arXiv preprint arXiv:1610.05492, 2016.

[74] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on mobile edge computing: The communication perspective," IEEE Communications Surveys & Tutorials, vol. 19, no. 4, pp. 2322-2358, 2017.

[75] S. B. Kotsiantis, I. D. Zaharakis, and P. E. Pintelas, "Machine learning: a review of classification and combining techniques," Artificial Intelligence Review, vol. 26, no. 3, pp. 159-190, 2006.

[76] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436-444, 2015.

[77] S. M. Erfani, S. Rajasegarar, S. Karunasekera, and C. Leckie, "High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning," Pattern Recognition, vol. 58, pp. 121-134, 2016.