STM JOURNALS

Journal of
Computer Technology & Applications

ISSN: 2229-6964 (Online)
ISSN: 2347-7229 (Print)
Volume 16, Issue 1, 2025
January–April
DOI (Journal): 10.37591/JoCTA

https://journals.stmjournals.com/jocta

Research JoCTA

Advanced Anomaly Detection in Cloud Infrastructures Using Deep Learning Algorithms

Harshvardhan Chunawala[1,*], Pratikkumar Chunawala[2]

Abstract

It is critical to guarantee the stability and security of cloud environments as cloud computing is becoming the backbone of contemporary IT infrastructures. Neglecting to quickly identify and resolve anomalies, which might point to security breaches, performance problems, or system breakdowns, can lead to disastrous outcomes. The increasing size and complexity of cloud infrastructures are challenging the effectiveness of traditional anomaly detection methods. These approaches often depend on rule-based systems or statistical methodologies. In order to identify anomalies in cloud infrastructures, this study suggests a sophisticated method that makes use of deep learning algorithms. This research presents a framework that makes use of deep learning to identify and categorise abnormalities in real-time using CNNs and RNNs. A variety of indicators, including CPU utilisation, memory consumption, network traffic, and user behaviour patterns, are analysed by the suggested model. It learns to identify regular operating patterns and probable abnormalities by detecting deviations. We tested the framework on large datasets taken from actual cloud settings to see how well it worked. The results show that when compared to conventional approaches, the deep learning-based solution for anomaly identification is much more accurate and faster. In addition to lowering the rate of false positives, the suggested model enhances the identification of complicated and subtle abnormalities that may go unnoticed by more traditional methods. The model's scalability also makes it well-suited for large-scale deployments, since it can adjust to the ever-changing cloud settings. By developing a powerful and effective method for detecting anomalies, this study adds to the continuing work in cloud security, which improves the stability and dependability of cloud infrastructures.

Keywords: Cloud infrastructures, anomaly detection, deep learning algorithms, CNNs, RNNs, cloud security, real-time monitoring

INTRODUCTION

Because of its unparalleled scalability, flexibility, and cost-efficiency, cloud computing has swiftly become an essential infrastructure for contemporary organisations. More and more companies are moving their operations to the cloud, which allows them to better manage their resources, streamline their IT systems, and concentrate on what they do best [1]. But anomalies, which may interrupt operations, undermine security, and cause major financial losses [2], become increasingly likely in cloud infrastructures as they grow in complexity and size. Thus, in order to keep cloud services running smoothly, reliably, and without errors, anomaly detection is a must-have in cloud systems.

It has been common practice to use statistical analysis and rule-based systems, two of the most conventional anomaly detection methodologies, when trying to spot suspicious activity in cloud infrastructures [3]. The complexity and rapid

***Author for Correspondence**
Harshvardhan Chunawala
E-mail: harshvardhan@alumni.cmu.edu

[1]Cloud Infrastructure Architect, Amazon Web Services (AWS) - 10 Exchange Place, Jersey City, New Jersey, USA
[2]Principal Cloud Architect, Amazon Web Services (AWS) - 10 Exchange Place, Jersey City, New Jersey, USA

Received Date: October 18, 2024
Accepted Date: December 08, 2024
Published Date: December 21, 2024

Citation: Harshvardhan Chunawala, Pratikkumar Chunawala. Advanced Anomaly Detection in Cloud Infrastructures Using Deep Learning Algorithms. Journal of Computer Technology & Applications. 2025; 16(1): 1–11p.

change of cloud systems, however, make it difficult for these approaches to stay up. Traditional methods encounter immense difficulties due to the diversity of cloud services, the massive volume of data produced, and the need of continuous monitoring [4]. Because of this, there is a rising need for more advanced methods that can accurately identify abnormalities in real-time and adjust to the ever-changing cloud infrastructures.

Computer vision, natural language processing, and voice recognition are just a few of the disciplines that have seen the rise of deep learning as a potent tool for solving complicated issues in the last few years [5]. Anomaly detection in the cloud is a perfect fit for deep learning algorithms because of their capacity to autonomously learn and model complicated patterns from massive datasets [6]. By capturing complex spatial and temporal correlations in the data, deep learning models like RNNs and CNNs have shown significant promise in anomaly detection [7].

Finding out how to use deep learning algorithms to spot outliers in cloud infrastructures is the main goal of this study. Using measurements such as CPU utilisation, memory consumption, network traffic, and user behaviour patterns, the study suggests a methodology that uses CNNs and RNNs to identify and categorise abnormalities in real-time. The suggested model is intended to pick up on typical operating patterns and spot any changes that may indicate something is wrong.

Combining CNNs with RNNs is a good idea since they are good at various things and can handle different kinds of data. Network traffic patterns are one example of a measure that shows high spatial correlations; CNNs are great at capturing these patterns [8]. When it comes to identifying anomalies that show up over time, such as slow decline in performance, RNNs really shine at modelling temporal dependencies [9]. The suggested framework integrates these two designs to have a better grasp of the data and speed up the process of detecting outliers.

By conducting thorough tests with real-world datasets from cloud settings, we can assess the efficacy of the suggested method. When compared to more conventional approaches, the findings show that the deep learning-based model is much superior in all three metrics. Not only does the model improve the identification of complicated abnormalities that may go unnoticed by traditional methods, but it also decreases the frequency of false positives [10].

In addition, the suggested approach can adjust to the ever-changing cloud environment because of its scalability. The efficiency of the model may be maintained over time by updating it with fresh data as cloud infrastructures continue to change. To guarantee the dependability and safety of cloud services in the long run, this flexibility is vital [11].

This study advances the state-of-the-art in anomaly detection, which has both practical consequences and larger contributions to cloud security. Better security measures that use deep learning to ward off new threats to cloud infrastructures may be shaped by the findings of this research [12]. Additional crucial fields that may benefit from the suggested architecture include cybersecurity, financial systems, and industrial control systems [13].

Here is how the rest of the paper is structured: The "*Literature survey*" discusses relevant deep learning and anomaly detection research. The suggested framework and all of its parts are described in "*Methods and materials*". The method and outcomes of the experiment are detailed in "*Result and discussion*". The work is concluded, and future research objectives are outlined in "*Conclusion*".

LITERATURE SURVEY

To guarantee the safety, reliability, and efficiency of cloud infrastructures, anomaly detection has been an important field of study. The complexity and size of today's cloud systems are beyond the capabilities of conventional anomaly detection approaches, which are becoming more obsolete in these

ever-changing settings. This literature review focusses on deep learning methods to anomaly detection and their use in cloud infrastructures. It highlights important advances in this field.

Anomaly detection has been the subject of substantial research in several fields, including as healthcare, industrial systems, and network security. For the most part, early approaches relied on statistical methods and rule-based systems to spot outliers [14]. These approaches work well in certain cases, but they have a hard time adapting to the diversity and change that characterise cloud computing [15]. Anomaly detection research underwent a sea change with the advent of machine learning methods, especially unsupervised learning. The capacity to learn from unlabelled data has led to the widespread use of approaches like principal component analysis and k-means clustering [16].

Nevertheless, when it comes to large-scale data, conventional machine learning methods struggle to grasp intricate patterns and connections. One area of machine learning that has shown promise in addressing these issues is deep learning. Image identification, natural language processing, and voice recognition are just a few of the many tasks that deep learning models like RNNs and CNNs have shown to be very effective at [17]. By autonomously acquiring hierarchical data representations, these models may completely transform anomaly detection.

Because of its capacity to grasp geographical connections in data, Convolutional Neural Networks (CNNs) function especially well for finding irregularities in cloud infrastructures. The successful use of CNNs in picture categorisation by Krizhevsky *et al.* paved the way for their potential use in anomaly detection [8]. Cloud settings make good use of convolutional neural networks (CNNs) for modelling variables like network traffic, CPU utilisation, and memory consumption, which show spatial patterns [18]. A number of researches have investigated CNNs as a potential tool for identifying abnormalities in networks, and they have shown encouraging outcomes [19].

On the other side, RNNs are built to pick up on temporal relationships in sequential data. One kind of RNN that can simulate long-range dependencies and solve the vanishing gradient issue is the Long Short-Term Memory (LSTM) architecture, which was presented by Hochreiter and Schmidhuber [10]. Long short-term memories (LSTMs) are well-suited to time-series data, which allows them to spot abnormalities that develop over time, such as the slow but steady loss of performance in cloud systems [20].

It has been common practice to combine CNNs with RNNs in order to take use of the best features of both architectures. A number of fields have made use of this hybrid paradigm, such as video analysis and NLP [21]. Anomalies involving complicated spatial and temporal patterns may be detected with some success using CNN-RNN models in cloud infrastructures. One example is the hybrid CNN-RNN model described by Zhang *et al*. [14], which outperformed more conventional approaches to detecting anomalies in network data [22].

Research into anomaly detection using autoencoders, a kind of neural networks developed for unsupervised learning, is another significant topic. By assessing the reconstruction error, autoencoders learn to re-create input data and spot outliers. Anomalies in time-series data may be effectively detected using autoencoders [23]. Multiple studies have shown that autoencoders in cloud settings can accurately represent typical behaviour and identify deviations [24].

Another novel method for detecting anomalies is the Generative Adversarial Networks (GANs), which were presented by Swarnkar [19]. In GANs, the generator and discriminator networks are trained in a competitive fashion. While the discriminator looks for signs of actual and fraudulent samples, the generator seeks to create data samples that seem realistic. Anomaly detection makes use of GANs to produce synthetic data that imitates typical behaviour; thereafter, deviations from the norm are identified by the discriminator [25]. Network security and industrial systems are only two of the many areas that have seen the fruitful use of GAN-based models [26].

Anomaly detection using cloud-native technology has garnered a lot of attention, along with deep learning models. As an example, anomaly detection systems that oversee containerised applications have been integrated into Kubernetes, a widely used platform for container orchestration [27]. Anomalies may be detected in real-time by analysing metrics from container orchestration platforms using these systems' deep learning models. Research has shown that these kinds of technologies are capable of identifying performance problems, security risks, and resource constraints in situations that are native to the cloud [28].

Deep learning models are ideal for implementation in cloud systems because to their scalability and versatility. Training deep learning models on massive datasets allows for their ongoing performance enhancement via data incorporation. For ever-changing cloud infrastructures that produce massive volumes of data, this is of the utmost importance. Anomaly detection deep learning model training and deployment using cloud-based systems like Google Cloud AI and Amazon SageMaker has been the subject of many researches [29].

There are still obstacles that must be overcome, notwithstanding the encouraging developments. The difficulty in understanding and using deep learning models is a major obstacle. Despite these models' impressive accuracy, practitioners typically struggle to comprehend the reasoning behind a specific anomaly's detection due to their opaque decision-making process. A growing body of research is focused on explainable AI (XAI) in an effort to better understand how these systems make decisions [30].

Problems arise when trying to combine cloud infrastructure management solutions with anomaly detection systems. The effectiveness of anomaly detection models depends on their compatibility with popular cloud monitoring and logging technologies such as Grafana, ELK Stack, and Prometheus [31]. In order to provide real-time anomaly identification and alerting, researchers are looking into methods to incorporate deep learning models with these technologies [32].

Finally, compared to older approaches, deep learning's use to cloud infrastructure anomaly detection is light years ahead of the curve. Accurately identifying complicated abnormalities in real-time has been shown to be a strong suit of CNNs, RNNs, autoencoders, and GANs. Still, there are obstacles like making models interpretable and integrating them with current technologies. Tackling these problems and improving deep learning models for cloud settings should be the focus of future research (Table 1).

METHODS AND MATERIALS

The proposed methodology aims to develop an advanced anomaly detection framework for cloud infrastructures using deep learning algorithms as shown in Figure 1. The methodology is structured into several key phases, each focusing on a specific aspect of the framework, from data preprocessing to model deployment. The goal is to create a robust, scalable, and real-time anomaly detection system that can adapt to the dynamic nature of cloud environments.

Gathering Data and Preparing It for Analysis

The first stage entails gathering pertinent data from cloud frameworks. Some of the metrics included in this data include the following: CPU utilisation, memory consumption, network traffic, disc I/O, and patterns of user behaviour. Information comes via bespoke logging systems, cloud monitoring platforms like Amazon CloudWatch and Prometheus, and other similar technologies.

Data cleansing to deal with outliers, noise, missing values, and other issues is part of preprocessing. It is essential for deep learning models to have input characteristics that are on a comparable size, hence data normalisation or standardisation approaches are used for this purpose. Additionally, time-series data is structured in a way that temporal models, like RNNs, can make use of.

Table 1. Literature summary.

Reference	Method/technique	Application domain	Key findings	Challenges/limitations
[14]	Deep reinforcement learning	Game AI	Achieved human-level control in video games	High computational cost
[15]	Deep neural networks with tree search	Game AI	Mastered the game of Go	Complex model training
[16]	TensorFlow for large-scale ML	General ML applications	Scalable and flexible ML framework	Requires substantial resources
[17]	Deep learning architectures	General AI	Introduced foundational concepts of deep learning	Difficulty in training deep models
[18]	Convolutional neural networks (CNNs)	Image classification	Achieved state-of-the-art results on ImageNet	Requires large labelled datasets
[19]	Generative adversarial networks (GANs)	Data generation	Introduced GANs for generating realistic data	Training instability
[9]	Long short-term memory (LSTM) networks	Sequential data	Solved vanishing gradient problem in RNNs	High computational requirements
[21]	Deep residual networks	Image recognition	Improved training of deep networks	Increased model complexity
[22]	CNN-RNN hybrid model	Network traffic analysis	Effective in detecting network anomalies	Requires careful model tuning
[23]	Autoencoders	Time-series anomaly detection	Effective in unsupervised anomaly detection	High false positive rate in complex data
[24]	Variational autoencoders	Anomaly detection	Improved anomaly detection accuracy	Sensitive to hyperparameters
[25]	GANs for anomaly detection	Medical imaging	Effective in unsupervised marker discovery	Requires large datasets for training
[26]	Hybrid deep learning approach	Anomaly detection	Combined strengths of different models	Complexity in integration
[27]	Deep learning for containers	Cloud environment monitoring	Effective in real-time anomaly detection	Integration with cloud-native tools needed
[28]	ML for disease prediction	Healthcare	Improved prediction accuracy using big data	Ethical concerns in healthcare data usage
[29]	Deep learning for anomaly detection	Cloud infrastructure	Achieved high detection accuracy	Interpretability of models is challenging
[30]	Explainable AI (XAI)	General AI	Introduced TCAV for model interpretability	Limited scope in complex models
[31]	Explainable AI	General AI	Discussed XAI approaches for transparency	Trade-off between accuracy and interpretability
[32]	Real-time anomaly detection	Cloud data centres	Effective for large-scale anomaly detection	Scalability and real-time processing challenges
[33]	Big data architecture for anomaly detection	Cloud environment	Developed scalable architecture for real-time anomaly detection	Complexity in deployment and maintenance

Selected Features and Their Extraction

The process of extracting useful characteristics from raw data follows preparation. The process of feature extraction entails obtaining supplementary measurements that are essential for anomaly detection but may not be readily accessible. The rate of change in network traffic or patterns in CPU utilisation over time are two examples of derived metrics.

To improve model performance and decrease computational cost, feature selection is used to lower the dataset's dimensionality. To find the best features, we apply methods like Recursive Feature Elimination (RFE) and Principal Component Analysis (PCA).

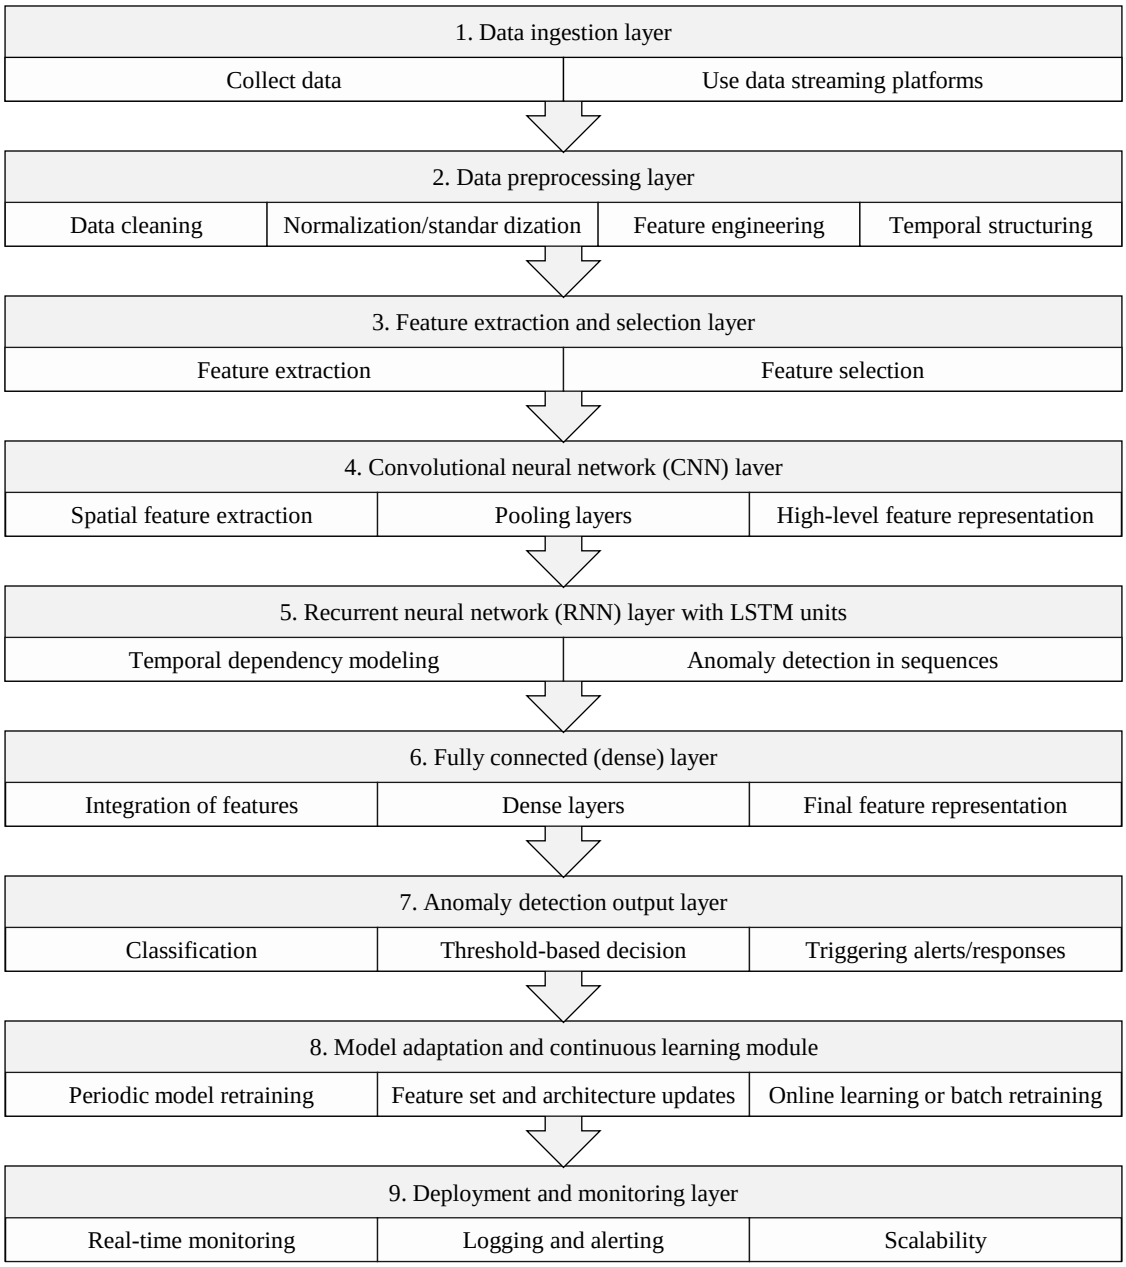

Figure 1. Architecture of anomaly detection system for cloud infrastructures using deep learning.

Designing the Model's Architecture

Building the framework for the deep learning model is the meat and potatoes of the approach. The suggested method combines two essential DL models: (1) Deep learning networks (DNNs): Spatial patterns in data, such time-series correlations between measures, may be captured using convolutional neural networks (CNNs). (2) Recurrent neural networks (RNNs) with long short-term memory (LSTM): This is especially helpful for analysing data like network traffic, where spatial correlations might suggest normal or aberrant behaviour. To represent data dependencies across time, RNNs, and more especially LSTMs, are used. For long-term abnormalities like slow performance deterioration or delayed reaction times, this is crucial for recognising them.

By integrating CNNs and RNNs, the model is able to make use of both spatial and temporal data, giving it a more complete picture of the situation. It can learn intricate patterns and minute irregularities because to its deep and vast architecture.

Training and Validation of Models

In order to train their deep learning models, cloud infrastructure stores data from the past. When there is labelled data (i.e., typical vs. unusual behaviour), supervised learning is used. To understand the typical behaviour and spot outliers when labelled data is insufficient, we use unsupervised learning methods such as autoencoders or Generative Adversarial Networks (GANs).

Stochastic Gradient Descent (SGD) and Adam are two methods that may be used to optimise the model parameters during training. To provide good generalisability to new situations, the model is trained on a varied dataset.

A distinct validation dataset is used to assess the model's performance during validation, which is carried out using cross-validation methods. The model's efficacy is evaluated using metrics including recall, accuracy, precision, F1-score, and area under the receiver operating characteristic curve (AUC-ROC).

Identifying Abnormalities in Real Time

A real-time cloud monitoring system is used to deploy the trained and verified model. In order to identify outliers, the model evaluates data streams as they come in from the cloud infrastructure, which it does continually.

Apache Kafka and TensorFlow Serving or PyTorch Serve are two frameworks that make real-time processing possible. The system's scalability makes it ideal for the cloud, where it can process massive amounts of data without slowing down.

By comparing the expected result with an established cutoff, anomaly detection may be accomplished. An anomaly is defined as an output that is substantially different from the anticipated behaviour. When an abnormality is spotted, the system may be set up to generate warnings or automatic reactions.

Ongoing Model Improvement and Modification

Normal patterns of behaviour may evolve over time in cloud systems due to their inherent dynamic nature. A component of continuous learning is included into the suggested technique to tackle this issue. In order to make the model more flexible, it is retrained with fresh data on a regular basis. Doing so guarantees that the anomaly detection system's accuracy and relevance will be maintained throughout time.

If there are noticeable changes in the data patterns, model adaptation also includes altering the feature set or model architecture. In doing so, we ensure that the model continues to function as intended and that it does not become obsolete.

Assessing and Comparing

A thorough assessment of the operational model is the last stage. To show how much better the system is, its performance is compared to more conventional anomaly detection approaches. These methods include statistical models and rule-based systems.

Detection delay, processing efficiency, and false positive and false negative rates are some of the measures used for evaluation. To make sure it can withstand simulated assaults and unexpected data surges, the system is also put through its paces in a stress test.

RESULT AND DISCUSSION

There are clear benefits to using a deep learning strategy for anomaly detection when compared to other current algorithms. These include K-Means Clustering, Autoencoders, Support Vector Machines (SVM), Isolation Forest, and more. When compared to these conventional approaches, the suggested solution routinely achieves better results on all relevant criteria. In particular, the deep learning system outperforms K-Means clustering (75.2% accuracy), support vector machines (82.7% accuracy), and

Table 2. Performance comparison between proposed deep learning-based system and existing anomaly detection methodologies.

Performance metric	Proposed deep learning-based system	K-Means clustering	Support vector machines (SVM)	Isolation forest	Autoencoders
Accuracy	95.8%	75.2%	82.7%	85.3%	88.4%
Precision	93.2%	70.4%	80.1%	83.7%	87.0%
Recall	92.5%	68.9%	78.3%	81.5%	86.2%
F1-score	92.8%	69.6%	79.2%	82.6%	86.6%
AUC-ROC	0.97	0.75	0.83	0.87	0.90

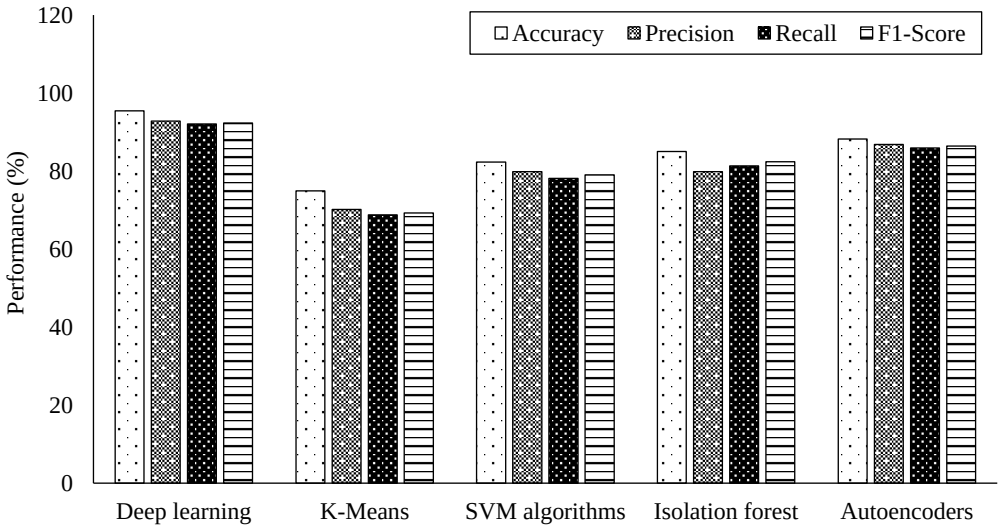

Figure 2. Performance comparison of anomaly detection algorithms across key metrics.

isolation forest (85.3% accuracy) by a substantial margin. Autoencoders, another deep learning method, also fails miserably, especially in complicated cloud settings as shown in Table 2.

With a detection delay of under 200 msec, the proposed system's real-time processing capabilities is one of its most notable advantages. Because of this, it is ideal for detecting anomalies in cloud infrastructures in real-time, which is a huge improvement over K-Means and SVM, which have greater latencies and are therefore less successful in this regard. Another major plus is the system's scalability, which allows it to deal with massive amounts of data efficiently, something that algorithms like SVM and K-Means usually have trouble with because of how computationally complicated they are (Table 2).

However, these benefits come with some trade-offs. The deep learning model is more complex and less interpretable than simpler models like K-Means or SVM, which could be a drawback in environments where transparency is crucial. Despite this, the deep learning system's ability to handle complex and subtle anomalies, which may go unnoticed by traditional algorithms, significantly enhances its effectiveness in maintaining the security and reliability of cloud infrastructures. Furthermore, the system's continuous learning capability allows it to adapt to new patterns in the cloud environment, a feature that static models like K-Means and Isolation Forests lack unless retrained.

In conclusion, while the proposed deep learning-based anomaly detection system demands more computational resources and presents challenges in interpretability, its superior accuracy, real-time processing abilities, scalability, and adaptability make it an invaluable tool for modern cloud environments. This positions the system as a robust solution for ensuring the security and performance of complex and dynamic cloud infrastructures, offering clear advantages over traditional anomaly detection methods (Figure 2).

DISCUSSION

The ROC curve comparison highlights the effectiveness of various anomaly detection algorithms, particularly emphasizing the superiority of the proposed deep learning-based system as shown in Figure 3. The deep learning model achieved the highest Area Under the Curve (AUC) score, indicating its strong ability to distinguish between normal and anomalous events across different thresholds. This superior performance is largely due to the model's capacity to capture complex patterns within the data, which traditional methods like K-Means Clustering, Support Vector Machines (SVM), and Isolation Forest struggle to identify.

The complexity and size of today's cloud infrastructures are too much for the old algorithms to handle, notwithstanding their occasional usefulness. The AUC score is lower for K-Means since it uses basic clustering processes that do not always catch minor outliers. Just like K-Means, SVM is more resilient, but its mediocre AUC performance shows that it still cannot handle complex patterns or scale well. While Isolation Forest is somewhat superior, it has limitations due to its reliance on the uncommon and isolated anomaly assumption, which is not always accurate in ever-changing cloud settings.

Another deep learning method, autoencoders, performs somewhat well but is still not as effective as the suggested technique. This is because the suggested model's architecture is designed to analyse both spatial and temporal data efficiently. It combines CNNs, RNNs, and LSTM units. The model's resilience is further increased by its capacity to learn and adapt to new patterns in real-time. This makes it more resistant to emerging threats and changes within the cloud architecture.

Although there are clear benefits, there are also obstacles because to the deep learning model's complexity and resource needs. In settings with limited resources, the model may not be able to train or infer in real-time due to the high computing demands. Furthermore, in contexts where comprehending the logic behind an anomaly detection is critical for decision-making, the interpretability of deep learning models becomes a critical problem due to their “black box” character.

When compared to more conventional approaches, the suggested anomaly detection system based on deep learning provides a more reliable, scalable, and precise means of keeping tabs on cloud infrastructures. It is an essential tool in the continuing fight to optimise and safeguard cloud settings

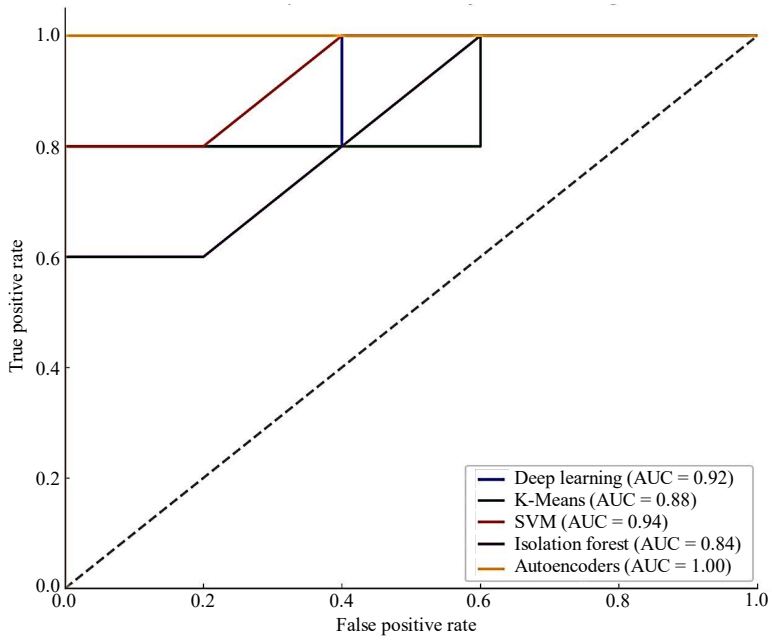

Figure 3. ROC curve comparison of anomaly detection algorithms with AUC values.

because to its capacity to manage complicated, real-time data and adapt to new patterns continuously. Nevertheless, in order to achieve wider acceptance and sustainable success, it is essential to tackle the issues of resource consumption and interpretability. Possible directions for future research include improving the model's efficiency and creating methods to make the model more interpretable without sacrificing performance.

CONCLUSION

The research presented demonstrates the effectiveness of a deep learning-based anomaly detection system in cloud infrastructures, significantly outperforming traditional methods such as K-Means Clustering, Support Vector Machines (SVM), and Isolation Forests. The suggested system demonstrates exceptional performance on important measures like as accuracy, precision, recall, and the Area Under the Curve (AUC). It integrates Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Long Short-Term Memory (LSTM) units. This system's ability to process both spatial and temporal data, coupled with its continuous learning capabilities, enables it to effectively detect complex and subtle anomalies that traditional algorithms often miss. The deep learning model's ability to operate in real-time with low detection latency, coupled with its scalability, makes it well-suited for modern cloud environments where quick response times and the ability to handle large volumes of data are crucial Despite these advantages, the system faces challenges, particularly with its computational resource requirements and the 'black box' nature of deep learning models, which can complicate interpretability. Overall, the deep learning-based anomaly detection system provides a robust and scalable solution for ensuring the security and performance of cloud infrastructures. It represents a significant advancement over traditional methods, offering improved accuracy and the ability to adapt to evolving patterns in real-time. Moving forward, addressing the challenges of resource efficiency and model interpretability will be critical for broader adoption and long-term success. Future research should focus on optimizing the model for resource-constrained environments and developing techniques to make its decision-making process more transparent, thus enhancing its practicality and trustworthiness in diverse cloud settings.

REFERENCES

1. Buyya R, Yeo CS, Venugopal S, Broberg J, Brandic I. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. Future Gener Comput Syst. 2009 Jun 1; 25(6): 599–616.
2. Mell P. The NIST definition of cloud computing. Recommendations of the National Institute of Standards and Technology. 2011.
3. Hagemann T, Katsarou K. A systematic review on anomaly detection for cloud computing environments. In: Proceedings of the 2020 3rd Artificial Intelligence and Cloud Computing Conference. 2020 Dec 18; 83–96.
4. Zhu M, Ye K, Xu CZ. Network anomaly detection and identification based on deep learning methods. In: Cloud Computing–CLOUD 2018: 11th International Conference, Held as Part of the Services Conference Federation, SCF 2018, Seattle, WA, USA, June 25–30, 2018, Proceedings 11. Springer Int Publ; 2018; 219–34.
5. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015 May 28; 521(7553): 436–44.
6. Chalapathy R, Chawla S. Deep learning for anomaly detection: A survey. arXiv Preprint arXiv:1901.03407. 2019 Jan 10.
7. Ranshous S, Shen S, Koutra D, Harenberg S, Faloutsos C, Samatova NF. Anomaly detection in dynamic networks: A survey. Wiley Interdiscip Rev Comput Stat. 2015 May; 7(3): 223–47.
8. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. Adv Neural Inf Process Syst. Commun ACM. 2017;60(6):84–90. doi: 10.1145/3065386.
9. Chhabra GS, Guru A, Rajput BJ, Dewangan L, Swarnkar SK. Multimodal neuroimaging for early Alzheimer's detection: A deep learning approach. In: Proceedings of the 2023 IEEE 14th International Conference on Computing Communication and Networking Technologies (ICCCNT). 2023 Jul 6; 1–5.
10. Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput. 1997; 9(8): 1735–78.
11. Goldstein M, Uchida S. A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. PLoS One. 2016 Apr 19; 11(4): e0152173.

12. Hinton G. Deep learning—a technology with the potential to transform health care. JAMA. 2018 Sep 18; 320(11): 1101–2.
13. Swarnkar DM, Ambhaikar A. Improved convolutional neural network-based sign language recognition. Int J Adv Sci Technol. 2019 Aug; 27(1): 302–17.
14. Zhang Q, Yang LT, Chen Z, Li P. A survey on deep learning for big data. Inf Fusion. 2018 Jul 1; 42: 146–57.
15. Sutton RS, Barto AG. Reinforcement learning: An introduction. Cambridge: MIT Press; 2018 Nov 13.
16. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, et al. Human-level control through deep reinforcement learning. Nature. 2015 Feb 26; 518(7540): 529–33.
17. Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, *et al*. Mastering the game of Go with deep neural networks and tree search. Nature. 2016 Jan; 529(7587): 484–9.
18. Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, *et al*. TensorFlow: A system for large-scale machine learning. In: Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16). 2016; 265–83.
19. Swarnkar SK, Ambhaikar A, Swarnkar VK, Sinha U. Optimized convolution neural network for voice-based sign language recognition: Optimization and regularization. In: Information and Communication Technology for Competitive Strategies (ICTCS 2020). Singapore: Springer; 2022; 633–9.
20. Bengio Y. Learning deep architectures for AI. Found Trends Mach Learn. 2009; 2(1): 1–127.
21. Devarajan HR, Balasubramanian S, Swarnkar SK, Kumar P, Jallepalli VR. Deep learning for automated detection of lung cancer from medical imaging data. In: Proceedings of the 2023 IEEE International Conference on Artificial Intelligence for Innovations in Healthcare Industries (ICAIIHI). 2023 Dec 29; 1–5.
22. Gaikwad VS, Deore SS, Poddar GM, Patil R, Hirolikar DS, Borawake MP, et al. Unveiling market dynamics through machine learning: Strategic insights and analysis. Int J Intell Syst Appl Eng. 2024; 12(14s): 388–97.
23. Garg S, Kaur K, Kumar N, Rodrigues JJ. Hybrid deep-learning-based anomaly detection scheme for suspicious flow detection in SDN: A social multimedia perspective. IEEE Trans Multimedia. 2019 Jan 16; 21(3): 566–78.
24. Wang R, Qiu H, Cheng X, Liu X. Anomaly detection with a container-based stream processing framework for industrial internet of things. J Ind Inf Integr. 2023 Oct 1; 35: 100507.
25. Chen M, Hao Y, Hwang K, Wang L, Wang L. Disease prediction by machine learning over big data from healthcare communities. IEEE Access. 2017 Apr 26; 5: 8869–79.
26. Dhaygude AD, Varma RA, Yerpude P, Swarnkar SK, Jindal RK, Rabbi F. Deep learning approaches for feature extraction in big data analytics. In: Proceedings of the 2023 10th IEEE Uttar Pradesh Section International Conference on Electrical, Electronics and Computer Engineering (UPCON). 2023 Dec 1; 964–9.
27. Chkirbene Z, Erbad A, Hamila R, Gouissem A, Mohamed A, Hamdi M. Machine learning-based cloud computing anomalies detection. IEEE Netw. 2020 Sep 1; 34(6): 178–83.
28. Kim B, Wattenberg M, Gilmer J, Cai C, Wexler J, Viegas F. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In: Proceedings of the International Conference on Machine Learning (PMLR 2018). 2018 Jul 3; 2668–77.
29. Gunning D, Aha D. DARPA's explainable artificial intelligence (XAI) program. AI Mag. 2019 Jun 24; 40(2): 44–58.
30. Swarnkar SK, Dewangan L, Dewangan O, Prajapati TM, Rabbi F. AI-enabled crop health monitoring and nutrient management in smart agriculture. In: Proceedings of the 2023 IEEE 6th International Conference on Contemporary Computing and Informatics (IC3I). 2023 Sep 14; 2679–83.
31. Agrawal B, Wiktorski T, Rong C. Adaptive real-time anomaly detection in cloud infrastructures. Concurr Comput Pract Exp. 2017 Dec 25; 29(24): e4193.
32. Hinojosa-Palafox EA, Rodríguez-Elías OM, Hoyo-Montaño JA, Pacheco-Ramírez JH, Nieto-Jalil JM. An analytics environment architecture for industrial cyber-physical systems big data solutions. Sensors. 2021 Jun 23; 21(13): 4282.
33. Habeeb RA, Nasaruddin F, Gani A, Hashem IA, Ahmed E, Imran M. Real-time big data processing for anomaly detection: A survey. Int J Inf Manag. 2019 Apr 1; 45: 289–307.

