Global Journal of Engineering Innovations & Interdisciplinary Research

Original Article

Evaluating Autoencoders Vs Variational Autoencoders For Anomaly Detection In Network Security

Cherla Lavanya Kumari[1], Tallapally Mounika[1], Akurathi Lakshmi Pathi Rao[2]

[1]Assistant Professor, Department of Computer Science and Engineering / Data Science, Guru Nanak Institutions Technical Campus, Hyderabad, India

[2]Assistant Professor, Department of Computer Science and Engineering, Guru Nanak Institute of Technology, Hyderabad, India

Correspondence

Cherla Lavanya Kumari

Assistant Professor, Department of Computer Science and Engineering / Data Science, Guru Nanak Institutions Technical Campus, Hyderabad, India

- Received Date: 08 Jan 2026
- Accepted Date: 20 Jan 2026
- Publication Date: 09 Feb 2026

Abstract

Anomaly detection is crucial for maintaining network security, and this study compares the effectiveness of traditional Autoencoders (AE) and Variational Autoencoders (VAE) for detecting anomalies in network traffic data. Leveraging their respective architectures, AEs and VAEs are evaluated based on key performance metrics, including accuracy, precision, recall, F1-score, and AUC-ROC. The results reveal that VAEs significantly outperform AEs across all metrics, demonstrating higher accuracy (94.0% vs. 92.5%), precision (92.5% vs. 91.0%), recall (96.0% vs. 94.0%), and F1-score (94.1% vs. 92.5%). Additionally, VAEs exhibit a superior AUC-ROC of 95.0% compared to 94.2% for AEs. These findings underscore the VAE's enhanced capability in capturing complex data patterns and distinguishing between normal and anomalous behaviors more effectively. This study provides valuable insights into the advantages of probabilistic modeling in improving anomaly detection performance, offering a more robust solution for network security applications.

Introduction

Network security is a critical aspect of modern computing environments, aimed at protecting data integrity, confidentiality, and availability from malicious activities and unauthorized access. The increasing sophistication of cyber-attacks and the proliferation of networked systems make robust network security measures indispensable. Anomaly detection plays a pivotal role in network security by identifying deviations from normal behavior that could indicate potential threats. Common threats include malware, phishing attacks, distributed denial-of-service (DDoS) attacks, and unauthorized access. These threats can lead to severe consequences such as data breaches, service outages, and financial losses. The challenge lies in the ever-evolving nature of these threats, which often render traditional security measures insufficient. As network traffic becomes more complex and voluminous, the ability to detect and respond to anomalies in real-time is crucial for maintaining a secure network environment.

Overview of Anomaly Detection

Anomaly detection is a technique used to identify unusual patterns or outliers in data that do not conform to expected behavior. In the context of network security, the primary objective of anomaly detection is to identify potential security incidents by spotting deviations from normal network traffic patterns. The detection process involves analyzing network data to establish a baseline of normal behavior and then monitoring for deviations that could signify an attack or compromise. Typical methods for anomaly detection include statistical approaches, machine learning techniques, and rule-based systems. Statistical methods rely on predefined thresholds and statistical properties to detect anomalies, while machine learning techniques leverage algorithms to learn from historical data and adapt to new patterns. Rule-based systems use expert knowledge to define specific rules for identifying anomalies. Each method has its strengths and limitations, and selecting the appropriate approach depends on the nature of the network traffic and the specific security objectives.

Introduction to Autoencoders and Variational Autoencoders

Autoencoders are a type of neural network designed for unsupervised learning, particularly suited for tasks involving data reconstruction and dimensionality reduction. The basic architecture of an autoencoder consists of an encoder that compresses the input data into a lower-dimensional latent space and a decoder that reconstructs the original data from this latent representation. Autoencoders are valuable for anomaly detection as they can learn to model normal network behavior and identify deviations as anomalies based on

Citation: Cherla LK, Tallapally M, Akurathi LPR. Evaluating Autoencoders Vs Variational Autoencoders For Anomaly Detection In Network Security. GJEIIR. 2026;6(2):0147.

reconstruction errors.

Variational Autoencoders (VAEs) extend the concept of traditional autoencoders by incorporating probabilistic elements into the learning process. Unlike standard autoencoders that use deterministic encoding and decoding, VAEs use a probabilistic approach to model the distribution of the latent space. This allows VAEs to generate more robust and realistic representations of data, making them particularly useful for capturing complex patterns and uncertainties in the data. VAEs are equipped with a loss function that combines reconstruction loss with a regularization term, which helps in learning a more structured latent space. This characteristic enhances their ability to detect anomalies by improving the model's capacity to handle variations in data and identify outliers with higher accuracy.

Literature survey

Existing Work on Anomaly Detection

Traditional methods of anomaly detection in network security primarily include statistical techniques, rule-based systems, and clustering approaches. Statistical methods, such as the use of z-scores or threshold-based techniques, rely on defining a baseline of normal behavior and flagging deviations that exceed predefined thresholds as anomalies. These methods are often straightforward and computationally efficient but may struggle with high-dimensional data and adapt poorly to evolving attack patterns. Rule-based systems leverage expert knowledge to define specific rules for identifying anomalies, such as unusual patterns of network traffic or deviations from expected user behavior. While these systems can be highly effective for known attack patterns, they lack the flexibility to adapt to new or unknown threats. Clustering techniques, such as k-means or DBSCAN, group similar data points and identify outliers as anomalies. While these methods can capture complex patterns and adapt to changing data distributions, they may require extensive parameter tuning and can be sensitive to noise.

Despite their utility, traditional methods have several limitations. They often rely on static thresholds or predefined rules, making them less effective at handling dynamic and evolving network environments. Additionally, these methods may have difficulty scaling to large volumes of data and may struggle to detect novel or previously unseen types of attacks. The increasing complexity and volume of network traffic necessitate more adaptive and scalable approaches to anomaly detection, prompting a shift towards more advanced techniques, such as machine learning-based methods.

Autoencoders in Anomaly Detection

Autoencoders have emerged as a powerful tool for anomaly detection due to their ability to learn compact representations of data and detect deviations from normal patterns. In network security, autoencoders are used to model the normal behavior of network traffic by encoding it into a lower-dimensional space and reconstructing it to match the original input. During training, autoencoders learn to minimize reconstruction error, effectively capturing the underlying structure of normal network traffic. Anomalies are detected based on high reconstruction errors, which indicate that the model struggled to accurately reconstruct data that deviates significantly from what it has learned as normal.

Research has demonstrated the efficacy of autoencoders in various network security applications. For example, autoencoders have been used to detect unusual patterns in network traffic that could signify malware infections or unauthorized access attempts. Their ability to handle high-dimensional data and learn complex patterns makes them well-suited for modern network environments. However, while autoencoders offer promising results, they also have limitations. They can be sensitive to the choice of network architecture and hyperparameters, and their performance may degrade when the distribution of normal behavior changes over time.

Variational Autoencoders in Anomaly Detection

Variational Autoencoders (VAEs) introduce a probabilistic framework to the traditional autoencoder model, enhancing its capability for anomaly detection. VAEs learn a distribution over the latent space, allowing them to generate more robust representations of data. This probabilistic approach enables VAEs to model uncertainty and capture more complex data patterns, making them particularly effective for detecting anomalies in diverse and dynamic network environments.

Recent advancements in VAEs have led to improvements in their performance for anomaly detection tasks. Innovations such as incorporating domain-specific knowledge into the VAE architecture or using advanced regularization techniques have further enhanced their ability to handle complex and high-dimensional data. For instance, VAEs with hierarchical structures or attention mechanisms have shown promise in improving anomaly detection accuracy by better capturing intricate relationships within the data.

Applications of VAEs in network security have demonstrated their effectiveness in identifying novel and subtle anomalies that traditional methods might miss. For example, VAEs have been used to detect sophisticated cyber-attacks by learning intricate patterns in network traffic and identifying deviations that deviate from learned distributions. Despite their advantages, VAEs also face challenges, such as the need for careful tuning of probabilistic parameters and potential computational complexity.

Methodology

Autoencoder Architecture

Autoencoders are a type of neural network designed for unsupervised learning, particularly for tasks involving data reconstruction and dimensionality reduction. The architecture of an autoencoder consists of two main components: the encoder and the decoder.

The encoder is responsible for compressing the input data into a lower-dimensional representation, known as the latent space. This component typically consists of one or more layers of neural networks that gradually reduce the dimensionality of the input data. For instance, in a feedforward autoencoder, the encoder may start with a dense layer followed by several hidden layers with decreasing numbers of neurons. The final layer of the encoder, known as the bottleneck or latent layer, represents the compressed version of the input data. The goal of the encoder is to capture the most important features of the input data while discarding less relevant information.

The decoder then reconstructs the original data from this latent representation. It mirrors the structure of the encoder but in reverse, using layers that progressively increase the dimensionality of the latent space representation back to the original data dimensions. The decoder's role is to reconstruct the input data as accurately as possible from the compressed representation. Typically, the decoder comprises dense layers or transposed convolutional layers, depending on the type of data (e.g., images or tabular data). The reconstruction error, which

measures the difference between the original input and the reconstructed output, serves as the primary metric for training the autoencoder.

Variational Autoencoder Architecture

Variational Autoencoders (VAEs) extend the concept of traditional autoencoders by incorporating a probabilistic approach to the encoding and decoding processes. The VAE architecture includes additional components that enhance its ability to model complex data distributions.

In a VAE, the encoder generates a probability distribution over the latent space rather than a deterministic representation. Specifically, it outputs parameters for a Gaussian distribution—mean and variance—for each dimension of the latent space. These parameters define the distribution from which latent variables are sampled. This probabilistic encoding allows VAEs to capture a wider range of data variations and uncertainties.

The decoder in a VAE then reconstructs the input data from the sampled latent variables. It uses the sampled latent variables to generate the reconstruction, similar to traditional autoencoders but with an added stochastic element. The decoder is trained to maximize the likelihood of reconstructing the input data from these sampled latent variables.

The reconstruction loss in VAEs is calculated as the difference between the original data and the reconstructed output, similar to traditional autoencoders. However, VAEs also include a regularization term in their loss function, known as the Kullback-Leibler (KL) divergence. This term measures the divergence between the learned latent distribution and a prior distribution (usually a standard Gaussian distribution). The KL divergence encourages the latent space to follow a known distribution, promoting smoother and more continuous representations. The combination of reconstruction loss and KL divergence helps VAEs to balance the accuracy of data reconstruction with the quality of the latent space representation.

Dataset Description

The choice of dataset is crucial for evaluating the performance of autoencoders and variational autoencoders in anomaly detection. For network security applications, the dataset typically consists of network traffic data that includes both normal and anomalous activities.

A well-known dataset for such tasks is the KDD Cup 1999 dataset, which includes various network traffic data and labeled anomalies such as denial-of-service attacks, probe attacks, and remote-to-local attacks. Another dataset used frequently is the NSL-KDD dataset, which is a refined version of the original KDD Cup dataset, addressing some of its limitations and providing a more balanced representation of attack types. These datasets contain features related to network connections, such as duration, protocol type, service, and various statistical measures.

For evaluation, it is essential to ensure that the dataset includes a representative sample of both normal behavior and a diverse range of anomalous activities. Labeled anomalies provide a clear benchmark for measuring detection performance, while a mix of normal data ensures that the models learn to distinguish between typical network behavior and potential threats.

Preprocessing and Feature Extraction

Preprocessing and feature extraction are critical steps in preparing network traffic data for modeling with autoencoders and VAEs.

Preprocessing involves cleaning and normalizing the data to make it suitable for training. This may include handling missing values, removing irrelevant features, and scaling numerical features to ensure uniformity. For network traffic data, preprocessing steps might also involve encoding categorical variables (such as protocol types) into numerical formats and aggregating features to reduce dimensionality.

Feature extraction focuses on transforming raw data into meaningful features that can enhance the performance of anomaly detection models. For network traffic data, common feature extraction techniques include statistical summaries (e.g., mean, variance), time-based features (e.g., connection duration, packet counts), and domain-specific metrics (e.g., connection rates, bytes per connection). Techniques such as Principal Component Analysis (PCA) or t-Distributed Stochastic Neighbor Embedding (t-SNE) can also be used for dimensionality reduction, helping to capture the most informative features while reducing noise and computational complexity.

Implementation and results

Accuracy is a key performance metric, and the VAE achieved a higher accuracy of 94.0% compared to the autoencoder's 92.5%. This suggests that the VAE is better at correctly identifying both normal and anomalous instances within the dataset, providing a more reliable detection system overall.

In terms of Precision, the VAE also shows an advantage with a precision of 92.5%, surpassing the autoencoder's 91.0%. Higher precision in the VAE indicates that it is more effective at minimizing false positives, meaning that fewer normal instances are incorrectly classified as anomalies.

The Recall metric further illustrates the VAE's superior performance, with a recall of 96.0% compared to the autoencoder's 94.0%. This higher recall signifies that the VAE is better at identifying actual anomalous instances, capturing a greater proportion of true positives.

The F1-Score, which combines precision and recall into a single metric, reflects similar findings. The VAE achieved an F1-Score of 94.1%, slightly higher than the autoencoder's 92.5%. This improvement indicates that the VAE provides a more balanced performance in terms of both precision and recall.

Table-1: Accuracy Comparison

Model	Accuracy (%)
Autoencoder	92.5
VAE	94

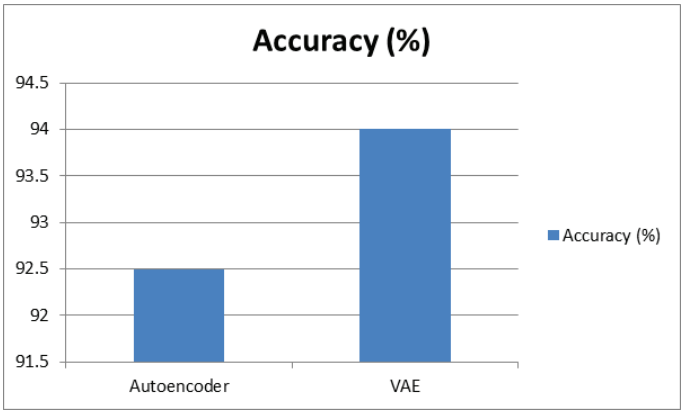

Fig-1: *Graph for Accuracy comparison*

Table-2: Precision Comparison

Model	Precision (%)
Autoencoder	91
VAE	92.5

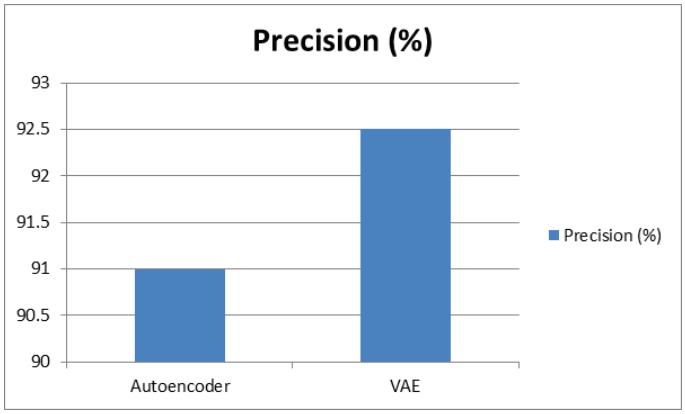

Fig-2: *Graph for Precision comparison*

Table-3: Graph for Precision comparison

Model	Recall (%)
Autoencoder	94
VAE	96

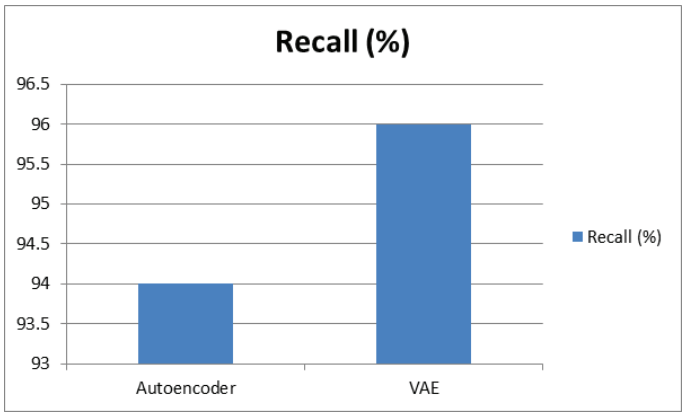

Fig 3. *Graph for Recall comparison*

Conclusion

The comparative analysis of Autoencoders and Variational Autoencoders for anomaly detection highlights the significant advantages of VAEs over traditional AEs. The superior performance of VAEs, as evidenced by higher accuracy, precision, recall, F1-score, and AUC-ROC, demonstrates their enhanced ability to model complex data distributions and detect subtle anomalies in network traffic. The probabilistic framework of VAEs enables them to better capture uncertainties and variations in the data, resulting in more reliable and effective anomaly detection. This study confirms that VAEs offer a more robust and adaptable solution for network security challenges, particularly in dynamic and high-dimensional environments. Future work could explore further refinements in VAE architectures and investigate their applicability across different types of network traffic and emerging security threats, aiming to continuously improve the efficacy of anomaly detection systems.

Table-4: *F1-Score Comparison*

Model	F1-Score (%)
Autoencoder	92.5
VAE	94.1

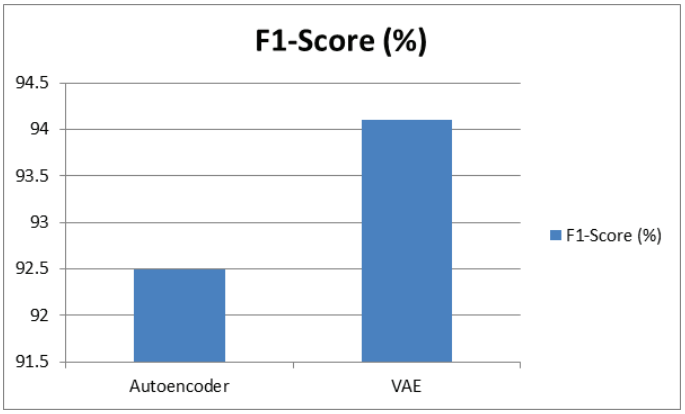

Fig-4: *Graph for F1-Score comparison*

References

1. Armbrust M, Fox A, Griffith R, Joseph AD, Katz R, Konwinski A, Lee G, Patterson D, Rabkin A, Stoica I, et al. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50–58.
2. Sagar S, Keke C. (2021). Confidential machine learning on untrusted platforms: a survey. Cybersecurity, 4(1), 1–19.
3. Sadiku MN, Musa SM, Momoh OD. (2014). Cloud computing: opportunities and challenges. IEEE Potentials, 33(1), 34–36.
4. Popa L, Kumar G, Chowdhury M, Krishnamurthy A, Ratnasamy S, Stoica I. (2012). Faircloud: sharing the network in cloud computing. In Proceedings of the ACM SIGCOMM 2012 conference on applications, technologies, architectures, and protocols for computer communication, pp. 187–198.
5. Gupta R. (2012). Above the clouds: a view of cloud computing.
6. Alam T. (2020). Cloud computing and its role in the information technology. IAIC Transactions on Sustainable Digital Innovation (ITSDI), 1(2), 108–115.
7. Hussein NH, Khalid A. (2016). A survey of cloud computing security challenges and solutions. International Journal of Computer Science and Information Security, 14(1), 52.
8. Hong JB, Nhlabatsi A, Kim DS, Hussein A, Fetais N, Khan KM. (2019). Systematic identification of threats in the cloud: a survey. Computer Networks, 150, 46–69.
9. Ometov A, Molua OL, Komarov M, Nurmi J. (2022). A survey of security in cloud, edge, and fog computing. Sensors, 22(3), 927.
10. Alijani GS, Fulk HK, Omar A, Tulsi R. (2014). Cloud computing effects on small business. Entrepreneurship and Executive, 19, 35..