The ISC Int'l Journal of Information Security

ISeCure

October 2025, Volume 17, Number 3 (pp. 1–11)

http://www.isecure-journal.org

Presented at the ISCISC'2025 in TEHRAN, Iran.

Lateral Movement Attack Detection using Variational Autoencoders **

Mostafa Shabani [1], and Tala Tafazzoli [2,*]

[1] *Department of Industrial Engineering, Iran University of Science and Technology*

[2] *ICT Security Faculty, ICT Research Institute (ITRC)*

ARTICLE INFO.

Keywords:
Lateral Movement Attack, Variational Auto Encoder, hybrid learning, Anomaly Detection

Type:

doi:

ABSTRACT

Lateral movement, a sophisticated cyberattack strategy, enables adversaries to stealthily infiltrate networks following an initial breach. Detecting such maneuvers is exceptionally challenging, as they are designed to seamlessly blend with legitimate system operations and network traffic, rendering traditional signature-based defenses ineffective. Supervised machine learning approaches, while promising, are constrained by their dependence on pre-labeled datasets of known attack patterns. To overcome these limitations, this study introduces a novel hybrid deep learning framework that integrates a Variational Autoencoder (VAE) for robust feature extraction, coupled with a supervised classifier to identify lateral movement. Through meticulous feature engineering on the LMD dataset, the VAE is trained exclusively on normative system and network behavior, constructing a probabilistic representation of legitimate activity. Anomalies, detected via reconstruction error, signal potential malicious intrusions. Empirical evaluation demonstrates the framework's superior performance, achieving a detection time of 00:00:02:54 and an AUC of 99.6983%, reflecting exceptional class separation and computational efficiency. This hybrid architecture delivers a scalable, high-accuracy solution, establishing the VAE as a pivotal tool for combating advanced persistent threats with unparalleled precision and operational viability.

1 Introduction

The clandestine propagation of malicious actors within compromised network perimeters, a process formally designated as lateral movement, constitutes a critical and often protracted phase in sophisticated cyber-assaults, particularly those orchestrated by Advanced Persistent Threat (APT) groups. Subsequent to an initial breach, these adversaries operationalize a diverse array of tactics, techniques, and procedures (TTPs) to traverse the internal infrastructure. The ultimate objective is typically the expropri-

* Corresponding author.

**The ISCISC'2025 program committee effort is highly acknowledged for reviewing this paper.

Email addresses: mostafa.shabanii@yahoo.com , tafazoli@itrc.ac.ir

ISSN: 2008-2045

ation of high-value assets, culminating in large-scale data exfiltration or systemic service disruption. The insidious nature of lateral movement lies in its capacity to obfuscate malicious activities by masquerading as legitimate network traffic, thereby rendering conventional signature-based and rule-based security paradigms largely ineffectual.

In response to this escalating threat landscape, machine learning methodologies have become indispensable. Nevertheless, extant approaches are circumscribed by inherent constraints. Purely supervised (discriminative) models, despite their high fidelity in classifying known threats, exhibit brittleness when confronted with novel attack vectors, a consequence of their fundamental dependence on pre-labeled training corpora. Conversely, purely unsupervised models, while proficient in detecting statistical anomalies indicative of zero-day threats, often fail to bridge the semantic gap required to categorize specific attack typologies and are susceptible to elevated false positive rates.

To transcend these limitations, this paper proposes and elucidates a novel semi-supervised framework predicated upon a Hybrid Variational Autoencoder-Classifier (HVAE-C) architecture. This framework operationalizes a multi-task learning paradigm, wherein the model is concurrently optimized against two distinct objectives. The primary objective is an unsupervised, generative task, whereby the Variational Autoencoder (VAE) learns a low-dimensional, semantically rich latent manifold that captures the underlying data distribution of system behaviors. The secondary objective is a supervised, discriminative task, in which an integrated classification module leverages this learned manifold to perform fine-grained classification of lateral movement techniques. A symbiotic relationship is thereby established: the generative objective functions as a potent regularization mechanism, constraining the feature space to enhance the discriminative model's generalization capabilities and mitigate the risk of overfitting.

The contributions of this research are principally threefold: First, the formalization and implementation of a novel semi-supervised, multi-task architecture that synergistically integrates generative and discriminative learning for the complex challenge of lateral movement detection. Second, a comprehensive empirical validation demonstrating that the proposed model achieves an optimized trade-off between high classification fidelity (AUC of 99.70%) and computational tractability (detection time of 2.54 seconds). Third, an elucidation of the strategic utility of this hybrid paradigm, which concurrently harnesses the precision of supervised learning for known threats and the exploratory power of generative models for novel anomaly detection, thus addressing a salient deficiency in the current state-of-the-art.

The remainder of this paper is organized as follows : Section 2 reviews the relevant literature. Section 3 provides key definitions for lateral movement and VAEs. Section 4 details our proposed hybrid architecture and feature engineering process. Section 5 presents the mathematical formulation. Section 6 describes the experimental setup and results. Finally, Section 7 discusses our findings, and Section 8 concludes the paper.

2 Literature Review

The application of deep learning to network intrusion detection has evolved significantly, moving from traditional machine learning towards more complex architectures. Early research highlighted the effectiveness of combining unsupervised feature extraction using stacked or sparse autoencoders with classifiers like Random Forest or Support Vector Machines [1] demonstrating superior accuracy and lower false alarm rates on benchmark datasets like KDD Cup '99 and NSL-KDD [1, 2]. Overviews of the field confirm a clear trend towards deep learning models, such as Deep Belief Networks (DBNs), Autoencoders (AEs), and Recurrent Neural Networks (RNNs) [3] [4], due to their ability to handle high-dimensional data and their potential to detect novel zero-day attacks [3]. However, challenges remain, including high computational costs [3], the need for more realistic datasets like UNSW-NB15 and CIC-IDS2017 [4, 5], and the "black box" nature of deep learning models, which complicates interpretability [5]. Recent research has increasingly focused on specific, sophisticated threats like lateral movement, leveraging host-based data such as Windows Event Logs and Sysmon data [6] [7, 8]. Methodologies in this area often involve extracting behavioral features from authentication and process creation logs, then applying powerful classifiers like Random Forest or XGBoost to distinguish malicious activity from benign administrative behavior [6–8]. Given the scarcity of public datasets for this specific task, researchers have often resorted to creating custom, semi-synthetic datasets by simulating attacks in controlled environments [6, 7].

3 Definitions

3.1 Lateral Movement in CyberSecurity

Lateral movement is a critical post-exploitation phase where an adversary, having established an initial foothold, systematically expands control across an internal network. This process is initiated with internal reconnaissance to map the network architecture and

identify high-value assets. A pivotal action within this phase is privilege escalation, through which attackers harvest credential artifacts like password hashes or Kerberos tickets to gain elevated permissions, often targeting domain administrator accounts. Armed with these credentials, adversaries employ stealthy Tactics, Techniques, and Procedures (TTPs) that mimic legitimate administrative behavior [9].

These include credential-theft attacks such as Pass-the-Hash (PtH) and Pass-the-Ticket (PtT), alongside the abuse of trusted remote services (e.g., RDP, SMB) and built-in system utilities like PowerShell and Windows Management Instrumentation (WMI) to execute code and pivot to other systems [10]. The successful culmination of a lateral movement campaign can lead to severe consequences, including large-scale data exfiltration, significant financial losses, and profound reputational damage. A primary objective is often to establish persistence, ensuring sustained access even if the initial point of entry is remediated. Detecting such clandestine activity poses a formidable challenge for security teams, as the TTPs are intentionally designed to blend with benign administrative traffic, thereby bypassing traditional signature-based security controls. This inherent stealth underscores the necessity for advanced security solutions capable of performing behavioral analysis to identify subtle anomalies and Indicators of Compromise (IoCs) across the network, enabling the neutralization of threats before they achieve their strategic objectives [11].

3.2 Variational Autoencoders (VAEs)

Variational Autoencoders (VAEs) are generative models consisting of an encoder and a decoder. The encoder maps high-dimensional input data to a lower-dimensional probabilistic latent space, where each latent variable is described by a mean (μ) and standard deviation (σ) of a probability distribution, typically a Gaussian. This probabilistic latent space enables the generation of novel data samples that are statistically similar to the training data. The decoder reconstructs the input by sampling from the latent distribution. This architecture not only facilitates data compression but also allows for the generation of new data instances, which is particularly valuable in tasks such as anomaly detection, where outliers can be identified as deviations from the learned distribution [12]. The VAE is trained by maximizing the Evidence Lower Bound (ELBO), which serves as a lower bound on the log-likelihood of the data. The ELBO consists of two components: the reconstruction loss, which measures how accurately the decoder reconstructs the input data, and the Kullback-Leibler (KL) divergence, which regularizes the latent space by ensuring the posterior distribution over the latent variables $q_\phi(Z|X)$ remains close to a predefined prior distribution, commonly a standard multivariate Gaussian, $p(z)N(\tilde{0}, I)$. The reconstruction loss is defined as [13]:

$$E_{q_\phi(Z|X)}[\log p_\theta(x|z)]$$

while the KL divergence is given by:

$D_{KL}(q_\phi(Z|X)|p(Z)) = \frac{1}{2}\Sigma_{j=1}^{n_z}(1 + log(\sigma_{\phi,j}^2(x) - \mu_{\phi,j}^2(x) - \sigma_{\phi,j}^2(x),$

where n_z denotes the dimensionality of the latent space and $\mu_{\phi,j}(x), \sigma_{\phi,j}^2(x)$ are the mean and variance of the j-th latent dimension, respectively.
To enable efficient gradient-based optimization, the reparameterization trick is employed, which re-expresses the latent variable z as a deterministic transformation of the encoder's outputs and an auxiliary noise variable ϵ, sampled from a standard normal distribution:

$Z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon, \epsilon \sim N(0, I),$

where $\odot$ denotes element-wise multiplication. This reparameterization enables backpropagation through the stochastic sampling process, facilitating efficient training. The VAE's optimization thus balances the trade-off between accurate data reconstruction and latent space regularization, which is critical for the model's ability to detect anomalies by identifying data points that deviate from the learned latent distribution [14].

4 Our Proposed VAE Architecture for Lateral Movement Detection

The framework designed in this research for detecting lateral movement is an advanced hybrid architecture that synergistically combines the principles of unsupervised representation learning with supervised classification. This design is intentionally chosen to harness the distinct advantages of both methodologies, thereby creating a detection system that is not only robust in identifying known threats but also retains the potential for discovering novel ones. At its core, the framework is a single, end-to-end deep learning model featuring a shared encoder with two distinct output branches: a reconstruction head (the decoder) and a classification head. The primary role of the VAE's unsupervised pathway is not merely anomaly detection but to serve as a powerful regularizer for feature extraction. By being tasked with reconstructing the input, the encoder is compelled to learn a robust and generalizable latent representation of the data, which in turn benefits the supervised classification task. This head is composed of fully-connected (Dense) neural network layers, culminating in a SoftMax activation function. Its express function

is to take the rich, compressed feature vector from the VAE's latent space as input and map it to the predefined output classes. The operational data flow within the model is as follows:

(1) Input Layer: The model accepts a 44-dimensional feature vector derived from Sysmon logs.
(2) Encoder Pathway: The input vector is passed through the encoder network, which outputs the parameters—mean (μ) and log-variance ($log(\sigma_2)$)— of a posterior distribution $q_\phi(z|x)$.
(3) Latent Space Sampling: A 16-dimensional latent vector, z, is sampled from this distribution using the re-parameterization trick ($z = \mu + \sigma \odot \epsilon$).
    (a) Dual Output Heads: The latent vector z is simultaneously passed to two parallel heads: a. The b.Decoder (Reconstruction Head), which aims to reconstruct the original 44-dimensional input from z.
    (b) The Classifier (Classification Head), which maps z to a probability distribution over the three output classes ('Normal', 'EoRS', 'EoHT').

Our hybrid architecture is designed with a dual-purpose advantage. Its supervised classifier head is optimized to accurately identify known attack patterns from the LMD dataset. Concurrently, the unsupervised VAE core enables the detection of novel or unknown anomalies. This is accomplished by monitoring the latent space; incoming data that produces a latent representation significantly different from the learned distributions of known classes is flagged as a potential new threat. Therefore, our model is more than a standard classifier; it is a tailored, end-to-end architecture that combines the discriminative power of supervised learning for known threats with the generative, anomaly-detecting capability of a VAE for unknown ones.

4.1 Input Data and Feature Engineering

The overall efficacy and performance of a VAE-based lateral movement detection system fundamentally rely on the careful selection and engineering of features from appropriate data sources. This research utilizes the LMD-2023 dataset, which is, to our knowledge, the only benchmark corpus comprising Sysmon logs specifically for evaluating lateral movement detection methods. This dataset is generated from EVTX log files and provides a comprehensive and contemporary collection of system events. The LMD-2023 dataset consists of approximately 1.75 million log samples, initially described by 93 features. It incorporates traffic from a wide range of state-of-the-art lateral movement techniques, which are categorized into two main classes of attacks: 'Exploitation of Remote Services (EoRS)' and 'Exploitation of Hashing Techniques (EoHT)'. The attack samples include legacy techniques like EternalBlue and Pass-the-Hash (PtH), as well as more recent threats such as Log4Shell, Follina, and SMBGhost. For the evaluation context, the dataset is structured as a multiclass problem with three labels: 'Normal', 'EoRS', and 'EoHT'. A critical characteristic of LMD-2023 is that it is highly unbalanced, with normal traffic constituting about 92% of the samples, while the EoRS and EoHT attack classes represent approximately 6% and 2%, respectively. This imbalance reflects real-world scenarios and poses a significant challenge for machine learning models, making the LMD dataset a robust benchmark for this research. These logs are considered a critical resource due to their comprehensiveness and accuracy in recording suspicious and routine network activities [11].
To effectively prepare this raw and heterogeneous data for the VAE model, appropriate and insightful feature engineering is essential. This process involves several critical steps:

- **Feature Extraction:** The initial stage involves extracting crucial features from raw network data and log files. While raw network traffic and system logs contain extensive and complex information, not all of it is necessary for model learning. Therefore, features such as the number of sent and received packets, connection duration, traffic volume, protocol types used, system activities, and inter-process interactions are extracted. These features represent various network and system behaviors during lateral movement attacks.
    - In this research, 93 distinct features were initially extracted from Sysmon log data. Features with insufficient information or negligible impact on deep learning model performance were subsequently removed based on scientific criteria and data analysis. Ultimately, 15 key features were selected for use in the deep learning models.
    - Temporal information from Sysmon logs, such as event timestamps, is critical for identifying suspicious activity patterns. Attackers often perform activities during specific, less monitored hours, especially late at night or on weekends, providing valuable insights for anomaly detection.
    - Feature importance analysis using model coefficients (logistic regression) and Principal Component Analysis (PCA) was employed to identify the most impactful features. Features with high absolute coeffi-

cients were retained, while those close to zero (e.g., below 0.05) were discarded. PCA helped in dimensionality reduction and prioritizing features with the most variance. Key features like 'ProcessId', which indicates executable process patterns, proved crucial for detecting malicious processes. Less significant features, including some temporal variables with low variance, IPv6-related features, and EventIDs lacking discriminative power, were removed to reduce complexity and prevent overfitting.

- **Handling Invalid Values:** Sysmon data may contain invalid, empty, or null values that can hinder model performance. These were identified (e.g., "NaN", "Null", very small decimal values). Approximately 11.73% of the data (101,738 out of 867,672 samples) contained invalid values. Strategies employed included removing samples lacking meaningful information or excessive noise, and replacing missing values with zero.
- **Categorical Feature Encoding:** Some Sysmon features are categorical (non-numerical) and cannot be directly processed by deep learning models. Examples include Computer (host name), DestinationPortName (e.g., HTTP, DNS), SourceIsIpv6, and temporal features like SystemTime_year. These are converted to numerical representations using techniques like one-hot encoding, which creates a separate binary column for each possible categorical value, enhancing model processability and pattern recognition. Consequently, the initial 15 conceptually selected key features were expanded into a final 44-dimensional feature vector for the model's input.
- **Numerical Feature Normalization:** Continuous numerical variables need to be scaled to a standard range to ensure optimal and uniform use by the deep learning model. Disparities in scale can lead to unequal influence during the learning process. Min-Max normalization is applied to convert feature values to a standard range (typically between 0 and 1), enabling the deep learning model to better identify patterns regardless of data scale differences and prevent over-optimization of certain features.
- **Noise Reduction:** Network data and logs often contain significant noise that can negatively impact model accuracy. Noise can appear as unusual activities, log errors, or irrelevant information. Noise reduction techniques are employed to clean the data and remove unnecessary information. For example, data with unusual patterns or sudden errors are filtered to retain clean and reliable data for model training.
- **Dimensionality Reduction:** Due to the large volume of extracted features and data, some features may overlap in information or have less importance in attack detection. Methods like Principal Component Analysis (PCA) are used to reduce feature dimensions, retaining only the more critical and relevant features. This not only reduces the computational complexity of deep learning models but also improves their performance.
- **Data Augmentation:** To enhance the performance of deep learning models and prevent overfitting, data augmentation techniques are utilized. This involves generating synthetic data or modifying existing data (e.g., altering feature values or applying random noise) to expand the dataset. Data augmentation allows the model to be trained on a more diverse set of data, strengthening its ability to identify various patterns.

4.2 VAE Network Architecture

The hybrid VAE architecture was implemented with specific layers and hyperparameters to optimize performance on the Sysmon log data. For the sake of reproducibility, the detailed structure is provided below. The encoder network is designed to map the 44-dimensional input feature vector into a compressed probabilistic latent space. It is composed of:

- An initial **Dense** layer with 128 units and a ReLU activation function.
- A **BatchNormalization** layer to stabilize training, followed by a Dropout layer with a rate of 0.2 to prevent overfitting.
- A second **Dense** layer with 64 units and ReLU activation.
- Another **BatchNormalization** layer.
- The final encoder layers consist of two parallel Dense layers, each with 16 units, which output the mean (μ) and log-variance (log_var) of the latent distribution.

The latent space dimensionality was set to 16. A **Lambda** layer then performs the reparameterization trick to sample a latent vector z from this distribution. The decoder network is structured to reconstruct the original 44-dimensional input from the latent vector z. Its architecture is symmetrical to the encoder:

- A Dense layer with 64 units and **ReLU** activation.
- A **BatchNormalization** layer followed by a Dropout layer (rate of 0.2).
- A second **Dense** layer with 128 units and ReLU activation.

- Another **BatchNormalization** and Dropout layer (rate of 0.2).
- The final output layer is a Dense layer with 44 units and a **Sigmoid** activation function, which scales the reconstructed features to a range between 0 and 1.

The integrated classifier branch takes the 16-dimensional latent vector z as input. It consists of a **Dense** layer with 32 units (**ReLU** activation), followed by **BatchNormalization**, and a final **Dense** output layer with 3 units and a **SoftMax** activation function for the three-class classification.
The careful selection of specific VAE network layers, their detailed configuration (e.g., number of units, filter sizes), and the overall depth of both the encoder and the decoder are crucial architectural considerations. These choices significantly impact the model's ability to effectively learn normal behavior and accurately detect deviations indicative of lateral movement. The VAE is specifically optimized to uncover hidden movement patterns within the processed Sysmon log data, thereby facilitating the detection of lateral movement attacks.

4.3 Supervised Classifier Branch

The final threat identification is performed by a supervised classifier branch integrated into our architecture. This branch takes the compressed latent vector, generated by the VAE's encoder, as its input. This vector serves as a high-level feature representation of the original system and network activities. The classifier's role is to map this representation to one of the three predefined classes ('Normal', 'EoRS', 'EoHT'). The final layer utilizes a SoftMax activation function to output a probability distribution across these classes, enabling precise, multi-class threat detection based on the powerful features learned by the VAE core. Our architecture is shown in Figure 1.

5 Mathematical Formulation and Anomaly Detection with VAE

The Variational Autoencoder (VAE) employed in this research learns to model the probabilistic distribution of host-level event logs data to detect lateral movement. This process is fundamentally driven by its ability to distinguish between normal and anomalous behaviors through reconstruction.

5.1 Loss Function and Model Training

The model was trained using a composite loss function designed to optimize both the VAE's generative capabilities and the classifier's discriminative performance simultaneously. The total loss is a weighted

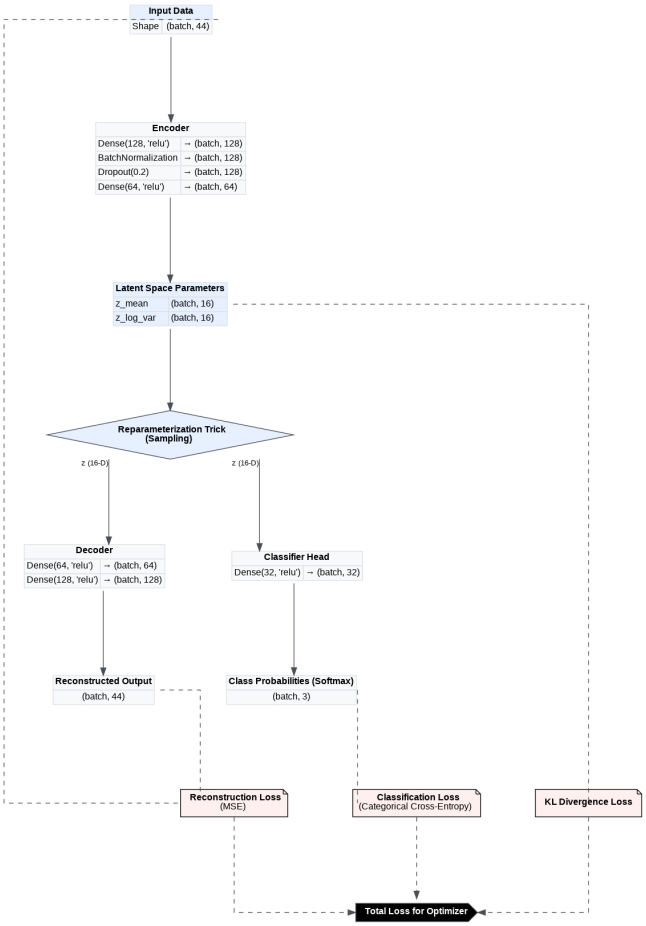

Figure 1. Our architecture for lateral movement attack detection

sum of three components:

(1) One **Reconstruction Loss:** This component quantifies how accurately the decoder reconstructs the original input data. It was implemented using the Mean Squared Error (MSE) between the input vector and the reconstructed output vector. This loss ensures the model learns a meaningful compression of the data.
(2) Two **Kullback-Leibler (KL) Divergence Loss:** This term acts as a regularizer on the latent space. It measures the divergence between the learned latent distribution (defined by z_{mean} and $z_{\log \text{var}}$ and a standard normal distribution. This encourages a continuous and well-structured latent space, which is critical for the model's generative properties.
(3) Three Classification Loss: This is the standard loss for the supervised part of the model. It was implemented using Categorical Cross-Entropy, which is ideal for multi-class classification problems. It measures the error between the predicted class probabilities and the true one-hot

encoded labels.

The overall training objective combined these losses. The reconstruction and KL divergence losses were added as a custom loss to the model, while the categorical cross-entropy was explicitly set as the loss for the classifier branch during model compilation. This hybrid approach allows the model to learn a robust feature representation while being fine-tuned for the specific task of lateral movement detection.

5.2 Anomaly Scoring and Detection

In our hybrid framework, the primary detection mechanism is the integrated supervised classifier. During training, the VAE's encoder learns to map high-dimensional input data to a robust latent representation, a process guided by both reconstruction loss and KL divergence. For the task of detection, this trained encoder acts as an intelligent feature extractor. An input sample is first encoded into its latent vector, which is then fed to the classifier branch. This branch makes the final prediction, assigning the sample to one of the 'Normal', 'EoRS', or 'EoHT' classes. Therefore, all performance metrics reported in this study are based on the output of this supervised classifier, not on a reconstruction error threshold. The higher the reconstruction error for a given data point, the more likely it is to be classified as an anomaly, signifying a deviation from the learned normal behavior. This mechanism allows the VAE to effectively identify patterns that are uncharacteristic of normal network operations and may signal malicious lateral movement activities.

6 Experimental Setup and Results

To rigorously evaluate the effectiveness and performance of the proposed VAE-based lateral movement detection method, a comprehensive experimental setup was established. This involved careful selection of appropriate datasets, thorough data preprocessing, and detailed model implementation and evaluation.

6.1 Dataset

This research primarily utilized the LMD (Lateral Movement Dataset) as the core data source for evaluating and benchmarking deep learning algorithms in lateral movement detection. The LMD dataset is a reference dataset specifically designed for assessing lateral movement detection methods. It comprises Sysmon logs generated from EVTX files, recognized as a primary and highly accurate source for studying lateral movement due to its comprehensiveness in recording suspicious and normal system and network activities. This dataset provides detailed information on attacker behavior during cyberattacks, making it a valuable tool for evaluating and comparing various algorithms. The LMD dataset is particularly suitable for deep learning methods because it contains real-world recorded events, especially those related to system processes and network activities, which are crucial for detecting complex patterns inherent in network traffic streams indicative of lateral movement.
The data collection methodology for the LMD dataset involved initially recording various system and network activities using advanced logging tools like Sysmon. Subsequently, EVTX (Event Log) files, containing detailed system event information, served as the primary input for generating this dataset. These files encompass various logs, including suspicious activities, network communications, file and process changes, and other critical system events that can reflect lateral movement patterns [15].
Our model was trained on the LMD-2023 dataset, derived from Sysmon event logs. We first engineered a 44-dimensional feature vector by refining an initial 93 features down to 15 key indicators, applying one-hot encoding to categorical data, and Min-Max scaling to numerical data.
To address the dataset's severe class imbalance (92% 'Normal' class), we employed down-sampling on the majority class. The resulting balanced corpus was shuffled and partitioned into stratified 80% training and 20% testing sets, with labels being one-hot encoded for the classification task.

6.2 Experimental Methodology

The proposed hybrid architecture for lateral movement detection was implemented using a deep learning framework, configured to effectively process the LMD dataset. The model's architecture, including the encoder and decoder components, was designed to learn complex patterns indicative of lateral movement. The model was trained using the Adam optimizer (learning rate: 0.001, batch size: 128) with a composite loss function that combined categorical cross-entropy for the classifier with the VAE's reconstruction (MSE) and KL-divergence losses. To prevent overfitting, EarlyStopping and ReduceLROnPlateau callbacks monitored the classifier's validation accuracy. Accordingly, the model's performance was evaluated using the classifier's metrics (e.g., Accuracy, F1-Score, AUC) rather than a reconstruction error threshold.
The models and experiments were implemented in Python 3.10.12 within the Google Colaboratory (Colab) cloud environment. We utilized the TensorFlow 2.15.0 framework and its integrated Keras library for building and training our model. Data preprocessing and manipulation were handled using the Pandas 2.0.3 and NumPy 1.25.2 libraries, while performance

metrics were computed using Scikit-learn 1.2.2. The training was accelerated on a Colab-provided runtime equipped with an NVIDIA Tesla T4 GPU containing 16GB of VRAM and 12.7 GB of system RAM to ensure computational efficiency. The environment was supported by CUDA Version 12.2.

6.3 Evaluation Metrics

To ensure a rigorous and quantitative assessment, the performance of the proposed lateral movement detection model was evaluated against a standard set of classification metrics. These included Accuracy, Precision, Recall, and the F1-Score to measure overall effectiveness and balance. Furthermore, the model's discriminative capability was assessed using the Area Under the Receiver Operating Characteristic Curve (AUC-ROC), while its error profile was analyzed through the False Positive Rate (FPR) and False Negative Rate (FNR).

6.4 Results

This section provides a meticulous exposition of the empirical performance demonstrated by the proposed hybrid architecture in the context of lateral movement detection. The evaluation is benchmarked against a diverse array of extant machine learning and deep learning methodologies, utilizing the LMD dataset as the empirical foundation. The hybrid model, architected to process 44 distinct features, was specifically configured for a three-class classification paradigm, encompassing 'Normal' activities, 'Exploitation of Remote Services' (EoRS), and 'Exploitation of Hashing Techniques' (EoHT). The quantitative assessment of model efficacy was conducted using a suite of standard performance indicators, namely: Area Under the Receiver Operating Characteristic Curve (AUC), Precision, Recall, F1-score, and Accuracy. Concurrently, considerations of computational overhead (Total Execution Time - T.E. Time) and the specifics of the training protocol (number of Epochs, application of k-fold cross-validation) were integral to the evaluation, as comprehensively detailed in Table 1.
The hybrid model demonstrates superior diagnostic performance, as evidenced by its ROC curve, which effectively illustrates the trade-off between the True Positive Rate (TPR) and False Positive Rate (FPR). The model achieves an impressive Area Under the Curve (AUC) of 99.6983%, signifying an exceptional ability to distinguish between normal and anomalous behaviors. While the VAE outperforms the LSTM network in terms of AUC (99.6983% vs. 95.82%), it shows slightly lower performance in terms of Precision (88.48% vs. 95.11%) and F1-score (88.50% vs. 95.55%). In contrast, the Extra Trees (ET) model exhibits the highest overall performance, surpassing the VAE in several metrics, including Accuracy (99.89%), Precision (99.05%), Recall (99.79%), and F1-score (99.41%).
While the Extra Trees (ET) model achieves a higher F1-score, our hybrid approach presents a critical trade-off that highlights its operational value. The hybrid framework offers a significant advantage that is paramount in real-world security operations: its ability to effectively handle high-dimensional data, enabling robust anomaly detection in complex systems.

(1) Exceptional Speed: With a total processing time of only 2.54 seconds compared to 7.5 minutes for the ET model, the model is exceptionally well-suited for real-time monitoring and rapid threat detection where latency is a major concern.

In summary, our hybrid model stands out as a highly efficient and effective tool for lateral movement detection. Its strengths in class separability, as demonstrated by the AUC, and its computational efficiency make it a valuable asset in cybersecurity applications. Although it has a lower macro-averaged F1-score compared to the ET and LSTM models, its capacity for hybrid learning and its potential to identify emerging threats highlight its operational viability. Future research should focus on enhancing recall and F1-score, particularly in multi-class scenarios, while preserving the model's established advantages in speed and discriminative power.

7 Discussion

The proposed hybrid approach for lateral movement detection is evaluated using key performance metrics like accuracy, precision, recall, F1-score, and AUC. These metrics demonstrate the model's ability to identify malicious activities while minimizing false positives. The Receiver Operating Characteristic (ROC) curve further illustrates its discriminative power across various thresholds. While the evaluation context is simplified, the VAE shows potential in high-stakes environments where detection failures are critical. However, the recall score of 88.55% in the 3-class evaluation suggests challenges in distinguishing between certain anomaly types, such as End of Reconnaissance (EoRS) vs. End of Honeypot Trigger (EoHT). This highlights an important operational trade-off: while a higher recall is desirable, the VAE's strength lies in its function as a high-speed, first-pass filter. In environments with massive data volumes, its ability to rapidly flag suspicious events for further analysis by human experts or more computationally intensive models is invaluable. The model prioritizes computational efficiency and broad anomaly detection over perfect classification of specific, known attack subtypes, which can be a strategic

choice in designing a layered defense architecture. The analysis also explores the influence of architectural choices and hyperparameter settings on performance. A comparison with other detection techniques, including graph-based methods and user behavior analysis, highlights our approach strengths in novelty detection and handling high-dimensional data. Nonetheless, the approach faces challenges, including sensitivity to training data quality, computational costs, and the complexity of interpreting the learned latent space. Future research could focus on refining VAE architectures, incorporating temporal data, improving latent space interpretability, and evaluating on diverse real-world datasets. Additionally, combining VAEs with methods like graph neural networks or Bayesian approaches could enhance detection capabilities.

8 Conclusion

This paper has presented a comprehensive exploration and evaluation of a hybrid architecture for the detection of lateral movement in network security. The detailed analysis of the VAE architecture and its application to anomaly detection has demonstrated its potential as a powerful tool for identifying the subtle activities associated with advanced persistent threats. The key contributions of this work are as follows:

- A Novel hybrid Framework: We proposed and successfully implemented a novel hybrid framework centered on a VAE, demonstrating the viability of generative models for this specific and complex cybersecurity challenge.
- A Compelling Balance of Accuracy and Efficiency: Our empirical results demonstrate that our architecture offers a strong balance between high discriminative power and operational efficiency. It achieves an excellent class-separation performance (AUC of 99.70%) while simultaneously providing superior computational speed, making it a viable solution for real-time security monitoring.

Continued research into optimizing VAE architectures, particularly through hybrid models, will be crucial for enhancing detection accuracy and robustness. Furthermore, addressing current limitations by employing parallelization to improve scalability and utilizing more diverse and noisy datasets will enhance the model's performance in operational environments. Ultimately, the VAE model, with its high potential in anomaly detection under hybrid conditions, stands as a powerful and efficient tool for combating complex cyberattacks. Further optimization and integration with other advanced techniques can make this model a vital instrument for cybersecurity defense against sophisticated threats.

While the current evaluation focuses on a supervised classification task, our hybrid architecture retains the potential for detecting zero-day threats. The reconstruction error from the VAE component can be monitored independently. A significant deviation from the norm in reconstruction error could signal a novel anomaly not seen during training. Exploring this unsupervised capability remains a promising direction for future research.

References

[1] Vu Dinh Phai Le Quy Don Qi Shi Nathan Shone, Tran Nguyen Ngoc. A deep learning approach to network intrusion detection. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2(1):41–50, 2018. .

[2] Mohammed Al-Habib Kamal Al-Sabahi Majjed Al-Qatf, Yu Lasheng. Deep learning approach combining sparse autoencoder with svm for network intrusion detection. *IEEE Access*, pages 52843–52856, 2018. .

[3] Jinoh Kim Sang C. Suh Ikkyun Kim Kuinam J. Kim Donghwoon Kwon, Hyunjoo Kim. A survey of deep learning-based network anomaly detection. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 22:949–961, 2019. .

[4] Cheah Wai Shiang Johari Abdullah Farhan Ahmad Zeeshan Ahmad, Adnan Shahid Khan. Deep learning approaches for anomaly-based intrusion detection systems: A survey, taxonomy, and open issues. *Transactions on emerging telecommunications technologies*, 31(1):949–961, 2021. .

[5] Ahmed Z. Emam Arwa Aldweesh, Abdelouahid Derhab. Deep learning approaches for anomaly-based intrusion detection systems: : A survey, taxonomy, and open issues. *Transactions on emerging telecommunications technologies*, 189(C), 2020.

[6] Mohammad A. Salahuddin Abbas Abou Daya Noura Limam Raouf Boutaba Tim Bai, Haibo Bian. Rdp-based lateral movement detection using machine learning. *Computer Communications*, 165(1):9–19, 2021.

[7] Konstantia Barbatsalou Christos Smiliotopoulos, Georgios Kambourakis. On the detection of lateral movement through supervised machine learning and an open-source tool to create turnkey datasets from sysmon logs. *International Journal of Information Security*, 22:1893–1919, 2023. .

[8] M. Roshni Thanka Ashwathy Anda Chacko, Bijolin Edwin. Detecting the lateral movement in cyberattack at the early stage using machine learning techniques. *Disruptive Technologies for*

Table 1. Comparative Performance Metrics for Lateral Movement Detection Models

Model	Features	Cis	AUC	Prec	Recall	F1	Acc	Epochs	K-Fold	T.E Time
MV (RF,LB,LoR) [16]	4	2	–	–	–	0.66	99.62%	N/A	✓	–
GRU DNN [13]	8	2	–	93.23%	–	–	96.68%	60	✓	–
Ensemble ML [17]	8	2	–	88.70%	–	–	–	N/A	✓	–
SS DL [18]	8	2	–	91.3%	–	–	99.9%	N/A	×	–
UML with JD [19]	15	2	–	6%	–	–	–	N/A	×	–
K-Means UML [20]	27	2	81%	–	–	–	–	N/A	×	–
RF [21]	29	2	–	83.73%	81.23%	0.82	–	N/A	✓	00:00:02:06
LaBi [22]	32	2	–	99.87%	99.47%	0.97	99.9%	N/A	✓	00:00:11:28
RF [23]	35	2	–	80.31%	80.29%	0.8	–	N/A	✓	00:00:03:11
ET [15]	15	3	99.84%	99.05%	99.79%	99.41%	99.89%	N/A	×	00:07:30:12
LSTM [15]	15	3	95.82%	95.11%	94.36%	95.55%	98.93%	30	×	00:15:44:18
VAE (Our method)	44	3	99.70%	88.48%	88.55%	88.50%	98.36%	58	✓	00:00:02:54

Big Data and Cloud Applications, 2021. .

[9] R. Zhang X. Wang, Z. Yan and P. Zhang. Attacks and defenses in user authentication systems: A survey. *Journal of Network and Computer Applications*, 188.

[10] D. Dimov and Y. Tzonev. Pass-the-hash. *Proceedings of the 18th International Conference on Computer Systems and Technologies*, 3.

[11] X. Chen S. Yu Q. Xuan J. Zhou, J. Yao and X. Yang. Lateral movement detection via time-aware subgraph classification on authentication logs.

[12] A. Kucukelbir D. M. Blei and J. D. McAuliffe. Variational inference: A review for statisticians. *J Am Stat Assoc*, 112(518):859–877, 2017.

[13] E. A. Barros da Silva L. Pinheiro Cinelli, M. Araújo Marins and S. Lima Netto. *Variational Autoencoder in Variational Methods for Machine Learning with Applications to Deep Networks*. Cham: Springer International Publishing, 2021.

[14] K. Dvijotham S. Gowal T. Cemgil, S. Ghaisas and P. Kohli. The autoencoding variational autoencoder. *Advances in Neural Information Processing Systems*.

[15] G. Kambourakis C. Smiliotopoulos and K. Barbatsalou. On the detection of lateral movement through supervised machine learning and an open-source tool to create turnkey datasets from sysmon logs. *Int J Inf Secur*, 22(6):1893–1919, 2023.

[16] G. Kaiafas et al. Detecting malicious authentication events trustfully. In *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*.

[17] G.-H. Syu C.-M. Chen and Z.-X. Cai. Analyzing system log based on machine learning model. *International Journal of Network Security*, 22(6).

[18] A. Fawaz A. Bohara, M. A. Noureddine and W. H. Sanders. An unsupervised multi-detector approach for identifying malicious lateral movement. In *Proceedings of the IEEE Symposium on Reliable Distributed Systems*.

[19] J. Liu B. Jiang L. Su M. Chen, Y. Yao and Z. Lu. A novel approach for identifying lateral movement attacks based on network embedding. In *IEEE Intl Conf on Parallel Distributed Processing with Applications, Ubiquitous Computing Communications, Big Data Cloud Computing, Social Computing Networking, Sustainable Computing Communications (ISPA/IUCC/BDCloud/SocialCom/SustainCom)*.

[20] A. Alva R. Sreedhar M. Bhadkamkar H. Pal Singh Bhasin, E. Ramsdell and H. Pal Singh. Data center application security: Lateral movement detection of malware using behavioral models.

[21] B. A. Powell. Role-based lateral movement detection with unsupervised learning. *Intelligent Systems with Applications*, 16.

[22] M. A. Salahuddin N. Limam A. A. Daya H. Bian, T. Bai and R. Boutaba. Uncovering lateral movement using authentication logs. *IEEE Transactions on Network and Service Management*, 18 (1):1049–1063, 2021.

[23] A. A. Daya M. A. Salahuddin N. Limam T. Bai, H. Bian and R. Boutaba. A machine learning approach for rdp-based lateral movement detection. In *IEEE 44th Conference on Local Computer Networks (LCN)*.

Mostafa Shabani holds a B.S. in Industrial Engineering from Semnan University and an M.S. in Financial Engineering from Iran University of Science and Technology. He is a specialist in Systems Analysis and Artificial Intelligence, whose professional focus is on applying machine learning algorithms for advanced modeling, strategic data analysis, and the development of intelligent decision-support systems.

Tala Tafazzoli received her PhD in Computer Engineering from AmirKabir University of Technology, Tehran, Iran. She is an assistant professor at ICT research institute (ITRC). Her research interests include blockchain technology, deep learning and cybersecurity. She has more than 25 years of research experience in e-commerce, and cyber security.