Large Language Model (LLM) Based Resilient Intelligent Transport Systems

Brij B. Gupta[1,2*], Akshat Gaurav[1], Valerie TANG[3], Varsha Arya[4], Amiya Nayak[5], Kwok Tai Chui[4]

[1]Department of Computer Science and Information Engineering, Asia University, Taichung 413, Taiwan
[2] VIZJA University, Warsaw, Poland
[3] Department of Supply Chain and Information Management, The Hang Seng University of Hong Kong, Shatin, Hong Kong
[4]Hong Kong Metropolitan University, Hong Kong SAR, China
[5] School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, ON K1N 6N5, Canada
bbgupta@asia.edu.tw, akshat.gaurav@ieee.org, tang9007@yahoo.com.hk, varshaarya2108@gmail.com, nayak@uottawa.ca, jktchui@hkmu.edu.hk

Abstract

Intelligent transport systems are increasingly dependent on interconnected devices and vehicular communications, making them vulnerable to reconnaissance attacks that precede large-scale intrusions. Traditional intrusion detection approaches often struggle with the scalability, redundancy of features, and complexity of dynamic traffic environments. To address these challenges, this paper introduces an LSH-Based Sparse Attention LLM Framework for reconnaissance attack detection. The framework applies Saint-Bowerbird Optimization (SBO) to select 23 optimal features from 41, ensuring efficiency and reducing redundancy. Categorical embeddings with drift analysis validate meaningful representation learning, while LSH-based sparse attention focuses on critical feature interactions with reduced complexity. The experimental results show that the model achieves 99.99% precision, outperforming RNN, LSTM, GRU and state-of-the-art models, confirming its robustness to secure intelligent transport systems.

Introduction

Intelligent transportation systems (ITS) are integral to the modern transportation infrastructure, connecting vehicles, road users, and traffic management systems, thus enhancing efficiency and safety. However, this interconnectedness exposes various vulnerabilities, especially to reconnaissance attacks, which involve gathering information to facilitate more severe types of cyberattacks. Effective detection mechanisms are required to protect ITSs against such threats.

An essential approach to defending against reconnaissance attacks within ITS involves the application of intelligent intrusion detection systems (IDS). These systems utilize advanced machine learning and deep learning methodologies, which have shown significant promise in identifying anomalous behavior indicative of reconnaissance activities. For example, the utilization of Convolutional Neural Networks (CNNs) has proven successful in detecting various network intrusions, particularly focusing on the subtleties of traffic patterns associated with reconnaissance efforts (Madhawa, Balakrishnan, and Arumugam 2018; Alarqan et al. 2025). This adaptability is crucial, as reconnaissance attacks often mimic legitimate traffic, making them challenging to identify with traditional methods.

The main contribution of this paper is an LSH-based Sparse Attention Large Language Model (LLM) framework for reconnaissance attack detection in intelligent transport systems. The use of the frozen LLaMA encoder is intended to leverage its strong contextual representation capability to model complex feature interactions in the tabular network data. Additionally, the LSH-based sparse attention mechanism is incorporated to improve feature dependency modeling and scalability, while maintaining computational efficiency in the proposed architecture.

The rest of the paper is organized as follows. Section 2 presents the details of the related work. Proposed model is explained in Section 3. Results are presented in Section 4. Finally, section 5 concludes the paper.

Related Work

Research in the domain of intelligent transport systems and IoT-enabled infrastructures has increasingly focused on intrusion detection, anomaly detection, and secure communication mechanisms. Several approaches have been proposed in recent years, each addressing different aspects of cybersecurity and resilience in smart vehicular environments.

Alkhudhayr and Ardah (2025) proposed an intrusion detection system (IDS) designed to mitigate cyberphysical risks in self-driving cars enabled by IoT. Their system integrated sensing devices and a controller area network simulator, where data were processed using a fine-tuned tree method driven by the search of krill herds. The model achieved high accuracy (98.4%) in identifying attacks, underscoring its effectiveness in improving the resilience of smart transport infrastructure. Similarly, Kadhim et al. (2025) introduced NOVA, a hybrid detection framework for vehicular networks that combined statistical anomaly detection with machine learning. The framework incorporated a trusted node mechanism with cryptographic authentication and demonstrated strong scalability in high-density environments, achieving 97% detection accuracy across multiple attack types. Zhou et al. (2022) proposed vehicle re-identification model.

*Corresponding Author

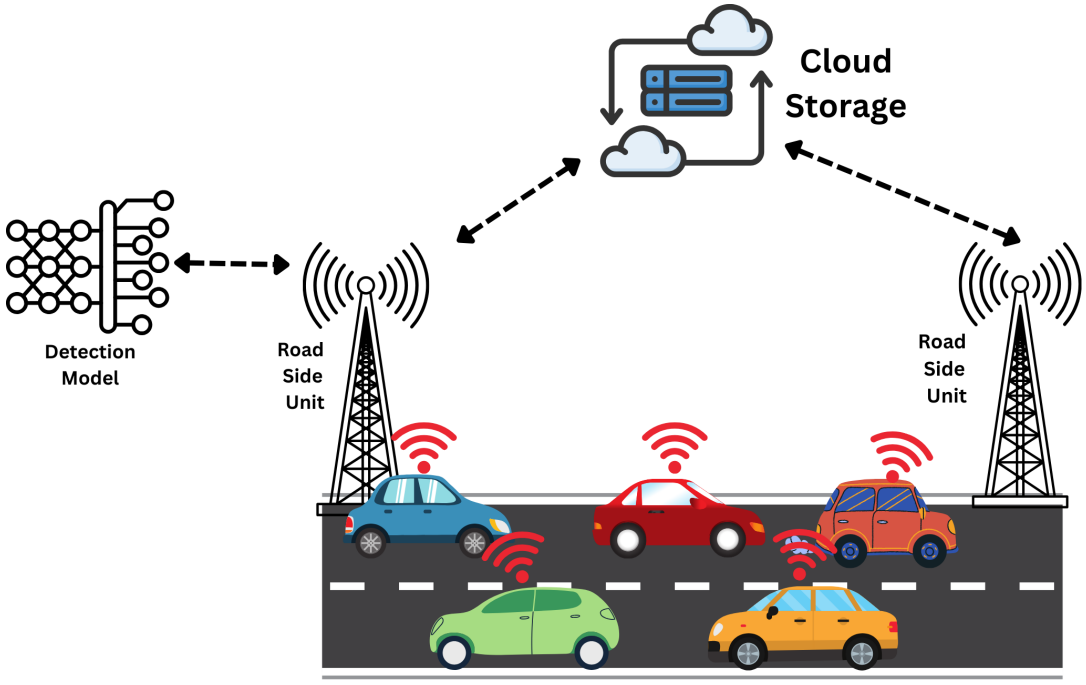

Figure 1: Proposed System Model

Deep learning has also been extensively explored to secure cyber-physical transport systems. AlEisa et al. (2023) investigated intelligent cyber–physical systems for safe vehicular communication using deep learning-based IDS. Their ensemble Long-Short-Term Memory (LSTM) model successfully detected cyber–physical attacks in both vehicle and external communication networks, achieving precision 99% on the UNSW-NB15 data set. Lilhore et al. (2024) contributed a different perspective by focusing on secure communication in the Internet of Vehicles (IoV). They proposed a hybrid encryption scheme that integrates AES-256, improved elliptic curve cryptography (IECC), and dynamic key management (DKM). Their approach reduced the transmission time by 30% compared to traditional AES-256 and achieved more than 99% intrusion prevention success, highlighting the critical role of cryptographic methods in resource-constrained vehicular environments. Gupta et al. (2022) proposed a model for secure smart vehicles.

Proposed Approach

System Model

The proposed LSH-Based Sparse Attention LLM Framework is installed in Roadside Units (RSUs) within the intelligent transport infrastructure, as represented in Figure 1. RSUs serve as critical intermediaries between vehicles and central servers, making them ideal for real-time traffic monitoring and decision making. By embedding the detection system directly into the RSUs, reconnaissance attacks can be identified close to the source, reducing latency and enabling faster defensive responses without relying solely on remote cloud servers.

The system monitors Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) traffic, which typically contains categorical and numerical features such as protocol type, flags, login attempts, and file access patterns. These traffic features are pre-processed, and Saint Bowerbird Optimization (SBO) is applied to select the most discriminative subset of attributes. The selected features are converted into categorical embeddings, whose stability is validated through drift analysis. The embeddings are then processed by the LSH-based sparse attention mechanism, which highlights the most relevant feature interactions while suppressing redundant computations. Finally, the LLM classifier learns the contextual dependencies between traffic flows, making a decision on whether the observed activity represents normal behavior or reconnaissance.

Detection Model

The detection model integrates categorical and continuous features to improve the detection of port scanning attacks. As shown in Figure 2, categorical variables are transformed into dense vector embeddings, while continuous variables are selected through BBO to retain only the most relevant attributes. Both feature types are combined and passed through a Linear Head Selection (LHS) attention mechanism. Unlike standard attention, the LHS attention mask restricts interactions to selected feature pairs, significantly reducing attention complexity while preserving discriminative dependencies across features.

The refined representation is then processed by an LLM that captures sequential dependencies and semantic correlations in the feature space. The LLM output is projected through fully connected linear layers to generate the final

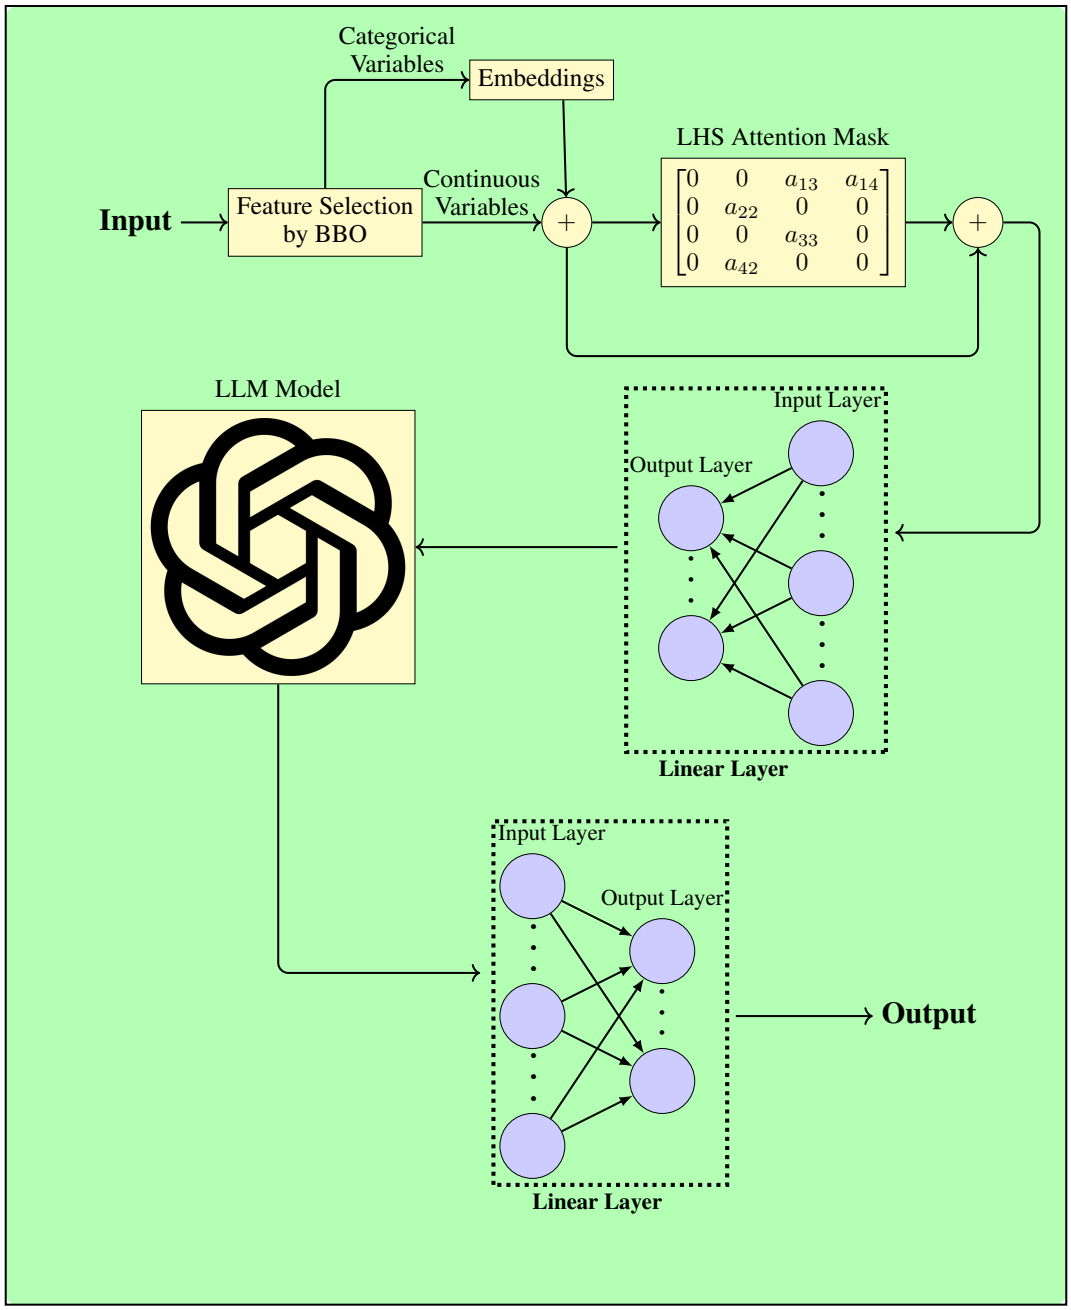

Figure 2: Detection Model

classification. This design enables efficient handling of high-dimensional inputs, reduces computational overhead due to LHS-based attention, and improves accuracy in distinguishing port scanning traffic from normal connections.

Raw Features

Let $x = (x^{\text{cat}}, x^{\text{num}})$ with C categorical fields and P numeric fields. Let the full index set be $\mathcal{I} = \{1, \ldots, C + P\}$.

BBO Feature Selection

Each habitat is a binary vector $z \in \{0, 1\}^{C+P}$ that encodes a subset $S(z) = \{i \in \mathcal{I} \mid z_i = 1\}$. The objective balances validation score and subset size

$$J(z) = \text{Score}\big(S(z)\big) - \lambda \, |S(z)| \tag{1}$$

where $\lambda > 0$. Migration and mutation update a population $\{z^{(m)}\}_{m=1}^{M}$ until convergence. The selected set is $\hat{S} = \arg\max_z J(z)$. We keep $x_{\hat{S}} = (x^{\text{cat}}_{\hat{S}_c}, x^{\text{num}}_{\hat{S}_p})$.

Embeddings and Assembly

For each kept categorical field $c \in \hat{S}_c$ use an embedding $e_c : \mathbb{Z} \to \mathbb{R}^{d_c}$. Stack them to $e(x) = \bigoplus_{c \in \hat{S}_c} e_c(x_c) \in \mathbb{R}^E$ with $E = \sum_{c \in \hat{S}_c} d_c$. Concatenate with selected numeric features to get

$$X = [\, e(x) \;\; x_{\hat{S}_p}^{\text{num}} \,] \in \mathbb{R}^{1 \times N}, \quad N = E + |\hat{S}_p|. \tag{2}$$

Feature Keys for LSH Attention

Learn keys $F \in \mathbb{R}^{N \times d}$ and use fixed projections $U \in \mathbb{R}^{h \times d}$. Hash codes

$$H = \mathbf{1}\{FU^\top > 0\} \in \{0,1\}^{N \times h}, \quad b_i = \sum_{t=1}^{h} 2^{t-1} H_{it}. \tag{3}$$

Neighbors $\mathcal{N}(i) = \{j \neq i \mid b_j = b_i\}$.

Pairwise Statistics

Distances $D_{ij} = \|F_i - F_j\|_2$. MI proxy $M_{ij} = \sigma(F_i^\top F_j)$. Two MLPs give α_{ij} and β_{ij}.

Sparse Attention with LSH Mask

$$A_{ij} = \begin{cases} \exp(-\alpha_{ij} D_{ij} + \beta_{ij} M_{ij}), & j \in \mathcal{N}(i), \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

$$\tilde{A}_{ij} = \frac{A_{ij}}{\sum_{k \in \mathcal{N}(i)} A_{ik} + \varepsilon}. \tag{5}$$

Aggregate features

$$\tilde{X} = X\, \tilde{A} \in \mathbb{R}^{1 \times N}. \tag{6}$$

Fuse original and attended

$$Z = [\, X \;\; \tilde{X} \,] \in \mathbb{R}^{1 \times 2N}. \tag{7}$$

Projection to LLM Space

Two linear layers with ReLU map to hidden size H

$$T = W_2\, \phi(W_1 Z^\top + b_1) + b_2 \in \mathbb{R}^{H \times 1}. \tag{8}$$

Transpose to a single token $T^\top \in \mathbb{R}^{1 \times H}$.

Frozen LLM encoder

Let g be the frozen LLaMA encoder.

$$h = g(T^\top) \in \mathbb{R}^{1 \times H}. \tag{9}$$

Classifier

Logits and probability

$$o = W_c h + b_c \in \mathbb{R}^{1 \times 2}, \tag{10}$$

$$\hat{y} = \mathrm{softmax}(o). \tag{11}$$

Cross entropy loss for label y

$$\mathcal{L} = -\sum_{c=1}^{2} y_c \log \hat{y}_c. \tag{12}$$

Complexity

Dense attention over N features needs $O(N^2)$ interactions. LSH limits each row to k neighbors with $k \approx N/2^h$ in expectation. Computation becomes $O(Nk)$ with hashing cost $O(Ndh)$. Thus the LHS mask yields a reduction from $O(N^2)$ to $O(Nk)$ with $k \ll N$.

Results and Discussion

Dataset Representation

To evaluate the proposed framework, a Kaggle data set (Zaib 2000) was utilized. The dataset contains two classes: Normal (67,342 samples) and Ipsweep (3,599 samples), as represented in Figure 3. Imbalanced data often bias the model toward the majority class, leading to poor detection performance for minority classes. To address this issue, a class weight balancing method was applied during training. This technique assigns higher weights to minority class samples and lower weights to majority class samples in the loss function. As a result, the model is penalized more for misclassifying minority class samples, which improves its sensitivity toward detecting reconnaissance attacks while maintaining overall performance.

Feature Selection by SBO

The Saint Bowerbird Optimization (SBO) algorithm was applied to select the most relevant subset of features from the original 41 features. After optimization, the algorithm identified 23 optimal features that contribute the most to the detection of reconnaissance attacks in intelligent transport systems. The performance of the SBO process is analyzed through multiple metrics and visualized in Figure 4.

Embedding Creation

To ensure that the categorical variables were correctly embedded in the model, the drift of the embedding across multiple categories was analyzed. The embedding drift plots in Figure 5 are used to visualize how the embeddings evolve in a reduced two-dimensional space during training. A stable drift indicates that embeddings are converging to meaningful positions, whereas unstable or overlapping drifts may indicate noisy or poorly separated embeddings.

LHS Based Sparse Attention

To efficiently capture dependencies among features while reducing computational cost, a locality-sensitive hashing (LSH)-based sparse attention mechanism was applied, as represented in Figure 6. Instead of computing dense pairwise attention across all features, LSH groups similar features into buckets and restricts attention computation within these buckets. This reduces quadratic complexity and focuses computation on the most relevant interactions.

Performance of LLM

The model was trained and tested on different data partitions by the use of a 70:30 train–test split. In order to reduce the randomness in the results, and assuring the reliability of

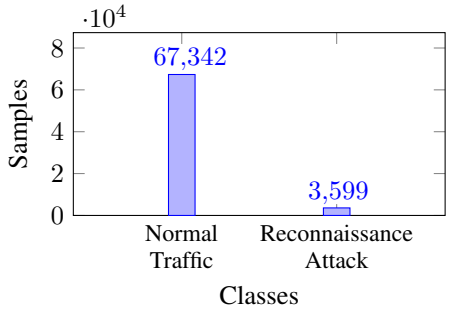

Figure 3: Class Distribution

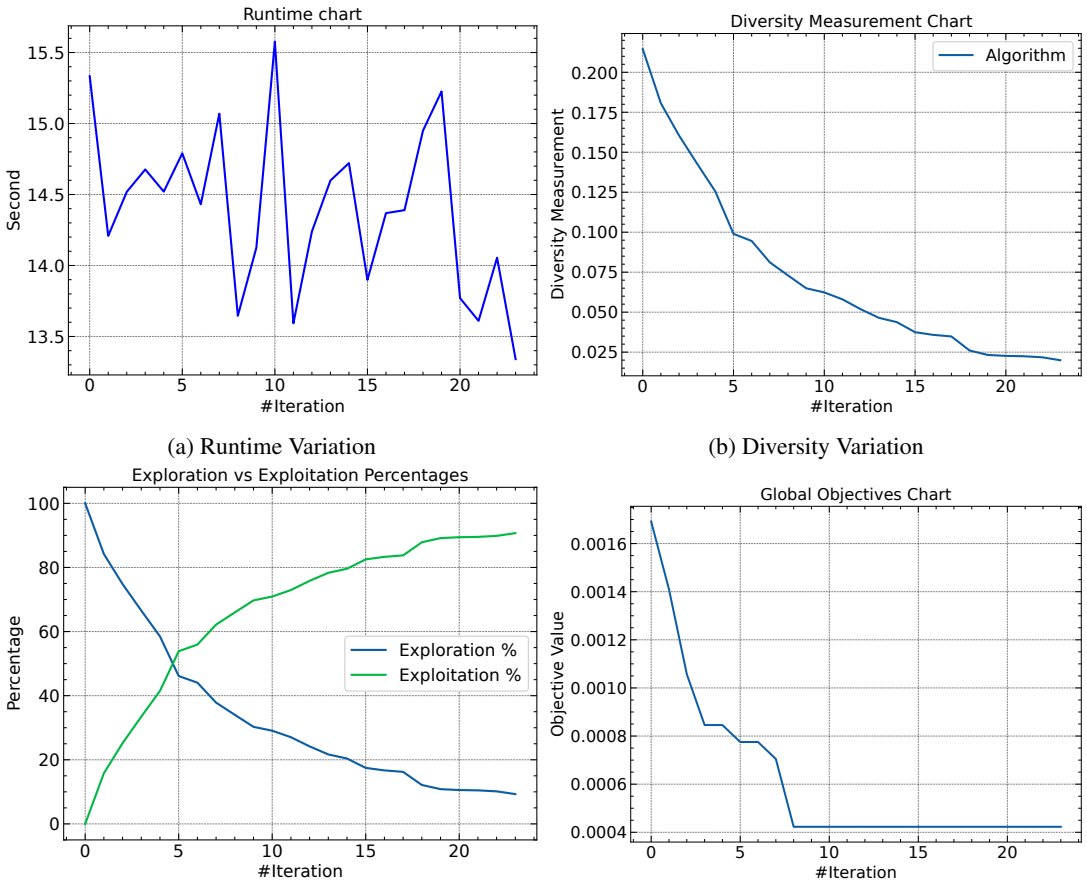

(a) Runtime Variation

(b) Diversity Variation

(c) Exploration versus Exploitation Variation

(d) Global Objective Variation

Figure 4: Performance of SBO

the experiments we performed a number of repeated experiments with fixed random seeds. The performance metrics reported are average results obtained through repeated experiments.

The Receiver Operating Characteristics (ROC) curve, the confusion matrix, and the classification report were used to plate the proposed LSH-based Sparse Attention LLM model. The evaluation measures, when considered together, offer a thorough insight into the model's proficiency in distinguishing normal traffic from reconnaissance traffic, along with its error distribution and precision-recall tradeoff.

The ROC curve, shown in Figure 7, demonstrates that the model achieves excellent discriminative capability. Both the Normal class and the Reconnaissance class reach an Area Under the Curve (AUC) score of 0.99984. The near-perfect AUC indicates that the model is highly effective in distinguishing between legitimate traffic attempts and reconnaissance attacks. The curve rises steeply toward the top left corner, reflecting both high sensitivity and high specificity.

The confusion matrix in Figure 8 provides detailed information on the classification outcomes. For the Normal class, 659 out of 662 samples were correctly identified, with

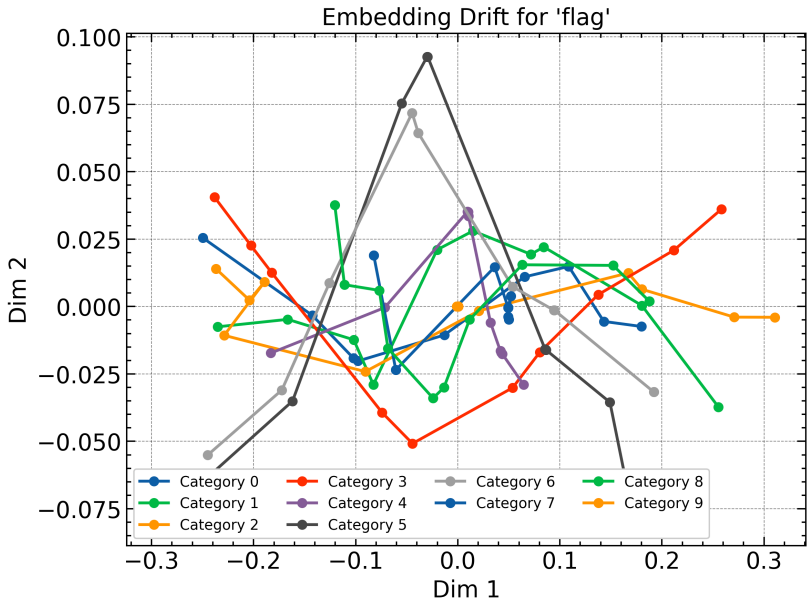

Figure 5: Embedding Drift

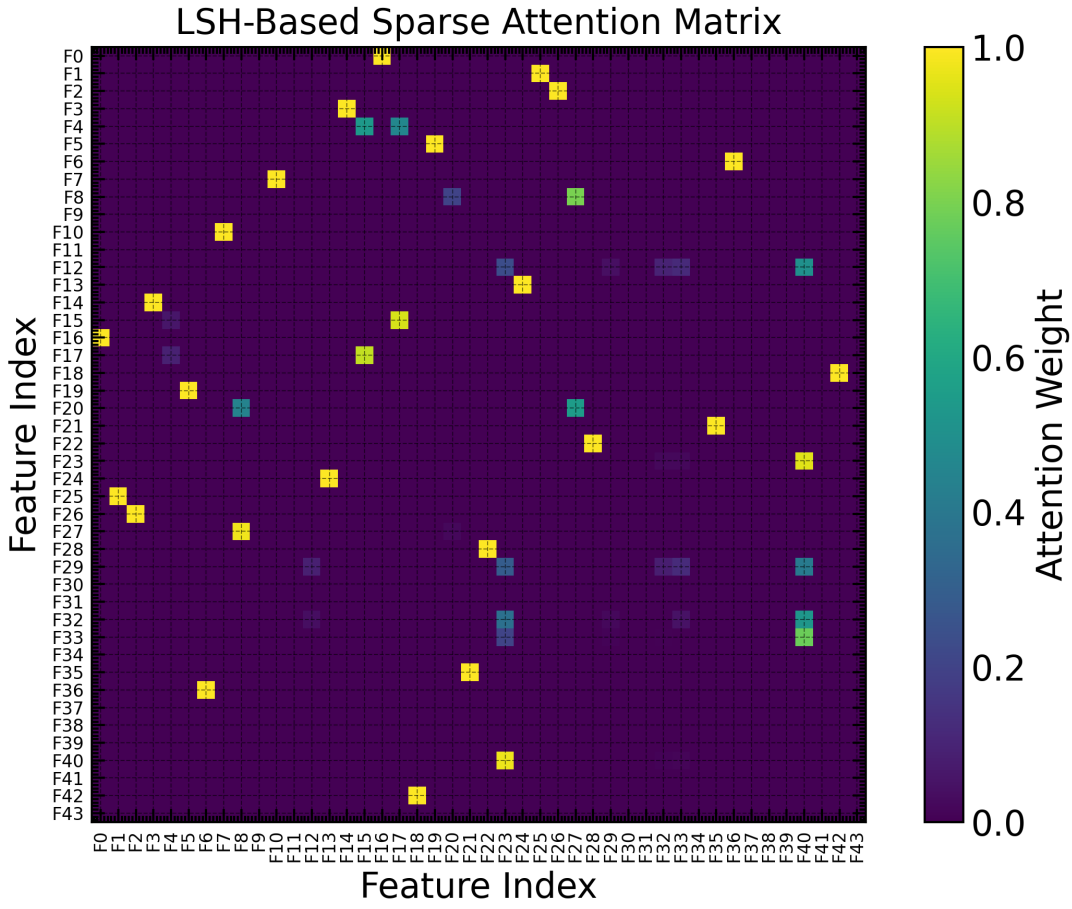

Figure 6: LHS Based Sparse Attention

only three instances misclassified as Reconnaissance. For the Ipsweep class, 13,360 samples were accurately detected,

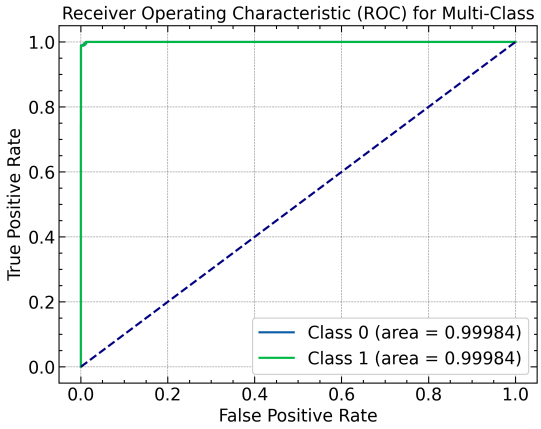

Figure 7: ROC

while 167 were mistakenly labeled Normal. This distribution shows that the model maintains very strong performance across both classes, though it exhibits a slight bias toward misclassifying a small fraction of attack samples as normal traffic.

Further analysis using the classification report in Figure 9 confirms this observation. For the Normal class, the model achieves a precision of 0.80, a perfect recall of 1.00, and an F1 score of 0.89, highlighting that while there are some false positives, the model does not miss legitimate samples. For the Reconnaissance class, precision reaches 1.00, recall is 0.99, and the F1 score is 0.99, reflecting almost flawless recognition of reconnaissance activity. The overall accuracy is recorded at 0.99, with the macro-average yielding a precision of 0.90, a recall of 0.99, and an F1 score of 0.94. The weighted average remains consistently high across all metrics, confirming that the performance of the model is not adversely affected by the inherent data imbalance.

Quantitative Comparison

To validate the effectiveness of the proposed LSH-based Sparse Attention LLM framework, its performance was compared to the baseline deep learning models, namely RNN, LSTM and GRU. The comparison focused on the

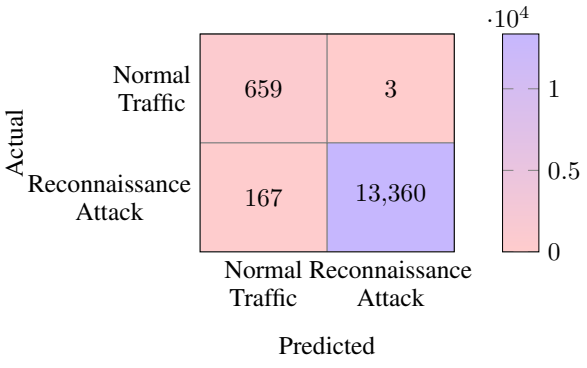

Figure 9: Classification Report

training loss across multiple epochs, as shown in Figure 10.

Comparison With State-of-the art models

To further evaluate the effectiveness of the proposed framework, a comparison was carried out against several recent state-of-the-art intrusion and attack detection models. Table 1 presents the results, highlighting the accuracy of different approaches along with the inclusion (or absence) of large language models (LLMs) and attention mechanisms.

Model	LLM	Attention	Accuracy
Alkhudhayr and Ardah (2025)	✗	✗	98.4
Kadhim et al. (2025)	✗	✗	97
AlEisa et al. (2023)	✗	✗	99
Lilhore et al. (2024)	✗	✗	99
Sharma, Babbar, and Sharma (2022)	✗	✗	70-98.2
Ashraf et al. (2021)	✗	✗	99.2
Mahamuni and Jalauddin (2021)	✗	✗	92
Proposed	LLama	LHS	99.99

Table 1: Comparison With State-of-the-Art Model

Figure 8: Confusion Matrix

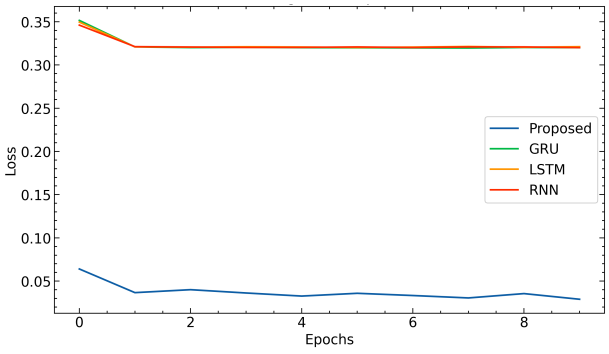

Figure 10: Loss Comparison

Conclusion Analysis

This paper proposes an LSH-Based Sparse Attention LLM Framework for Reconnaissance Attack Detection in Intelligent Transport Systems. Intelligent transport infrastructures are highly vulnerable to reconnaissance threats that exploit communication channels and system vulnerabilities. To address this challenge, the framework integrates Saint Bowerbird Optimization (SBO) for optimal feature selection, reducing 41 features to 23, and employs categorical embeddings with drift analysis to ensure stable representation learning. An LSH-based sparse attention mechanism is applied to capture the most relevant feature interactions while minimizing computational cost. The experimental evaluation demonstrates superior performance, achieving 99. 99% precision, with strong ROC, confusion matrix, and classification metrics. Comparisons with baseline (RNN, LSTM, GRU) and state-of-the-art models confirm the scalability and robustness of the framework for securing intelligent transport systems.

Acknowledgments

This research work is supported by National Science and Technology Council (NSTC), Taiwan Grant No. NSTC 114-2221-E-468-015

References

Alarqan, M.; Belaton, B.; Almomani, A.; Alauthman, M.; Al-Betar, M. A.; and Arya, V. 2025. Information theory-based ddos attack detection in cloud computing: A systematic survey of approaches, challenges, and future directions. *International Journal of Cloud Applications and Computing (IJCAC)*, 15(1): 1–38.

AlEisa, H. N.; Alrowais, F.; Allafi, R.; Almalki, N. S.; Faqih, R.; Marzouk, R.; Alnfiai, M. M.; Motwakel, A.; and Ibrahim, S. S. 2023. Transforming transportation: Safe and secure vehicular communication and anomaly detection with intelligent cyber–physical system and deep learning. *IEEE Transactions on Consumer Electronics*, 70(1): 1736–1746.

Alkhudhayr, H.; and Ardah, H. 2025. Mitigating cyberphysical risks in IoT-enabled smart transport infrastructure. *The Journal of Supercomputing*, 81(2): 446.

Ashraf, J.; Keshk, M.; Moustafa, N.; Abdel-Basset, M.; Khurshid, H.; Bakhshi, A. D.; and Mostafa, R. R. 2021. IoTBoT-IDS: A novel statistical learning-enabled botnet detection framework for protecting networks of smart cities. *Sustainable Cities and Society*, 72: 103041.

Gupta, B. B.; Gaurav, A.; Marín, E. C.; and Alhalabi, W. 2022. Novel graph-based machine learning technique to secure smart vehicles in intelligent transportation systems. *IEEE transactions on intelligent transportation systems*, 24(8): 8483–8491.

Kadhim, A. A. A. K.; Alzamili, Z. M.; Al-Shareeda, M. A.; and Amin, M. 2025. NOVA: A hybrid detection framework for misbehavior in vehicular networks.

Lilhore, U. K.; Simaiya, S.; Dalal, S.; Sharma, Y. K.; Tomar, S.; and Hashmi, A. 2024. Secure WSN architecture utilizing hybrid encryption with DKM to ensure consistent IoV communication. *Wireless Personal Communications*, 1–29.

Madhawa, S.; Balakrishnan, P.; and Arumugam, U. 2018. Data driven intrusion detection system for software defined networking enabled industrial internet of things. *Journal of Intelligent & Fuzzy Systems*, 34(3): 1289–1300.

Mahamuni, C. V.; and Jalauddin, Z. M. 2021. Intrusion monitoring in military surveillance applications using wireless sensor networks (WSNs) with deep learning for multiple object detection and tracking. In *2021 International conference on control, automation, power and signal processing (CAPS)*, 1–6. IEEE.

Sharma, A.; Babbar, H.; and Sharma, A. 2022. Ton-iot: Detection of attacks on internet of things in vehicular networks. In *2022 6th International Conference on Electronics, Communication and Aerospace Technology*, 539–545. IEEE.

Zaib, M. H. 2000. NSL-KDD.

Zhou, Z.; Li, Y.; Li, J.; Yu, K.; Kou, G.; Wang, M.; and Gupta, B. B. 2022. GAN-Siamese network for cross-domain vehicle re-identification in intelligent transport systems. *IEEE transactions on network science and engineering*, 10(5): 2779–2790.