

OPEN ACCESS

Received: 28/07/2026

Accepted: 21/09/2026

Published: 11/10/2026

Citation: Basil Raja P, Arockiam L, Aloysius A (2026) Cross-Domain Behavioural Mapping for IoT Attack Detection in Educational Networks Using CICIoT-2023 Dataset. Indian Journal of Science and Technology 19(35): 2461-2474. <https://doi.org/10.17485/IJST/v19i35.1057>

* **Corresponding author.**

basilraja1310@gmail.com

Funding: None

Competing Interests: None

Copyright: © 2026 Basil Raja et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Published By Indian Society for Education and Environment ([iSee](#))

ISSN

Print: 0974-6846

Electronic: 0974-5645

Cross-Domain Behavioural Mapping for IoT Attack Detection in Educational Networks Using CICIoT-2023 Dataset

P Basil Raja^{1*}, L Arockiam¹, A Aloysius¹

¹ Department of Computer Science, St. Joseph's College (Autonomous), Tiruchirappalli – 02, Affiliated to Bharathidasan University, Tamil Nadu, India

Abstract

Objectives: The study develops a cross-domain behavioural intrusion-detection framework for transferring attack knowledge from CICIoT-2023 to a derived educational IoT representation. **Method:** The proposed Cross-Domain Behavioural Mapping Intrusion Detection System (CDBM-IDS) integrates an Educational IoT Digital Twin Layer (E-DTL), Behaviour Spectrum Encoder (BSE), Hierarchical Deviance Encoder (HDE), Tri-Context Fusion (TCF), and Sector-Transfer Consistency Loss (STC-Loss). EduTwin is generated from CICIoT-2023 through educational device-role mapping, interaction modelling, and sector-oriented behavioural transformation. **Findings:** Under the controlled EduTwin evaluation protocol, CDBM-IDS achieves 98.72% accuracy, 98.64% precision, 98.51% recall, 98.57% F1-score, and 99.21% AUC-ROC. Ablation analysis confirms measurable contributions from BSE, HDE, TCF, and STC-Loss. **Novelty:** The framework introduces educational role abstraction, frequency-domain behavioural encoding, hierarchical deviation modelling, geometric context fusion, and consistency-driven sector mapping.

Keywords: IoT security; educational network; intrusion detection; cross-domain learning; behaviour modelling

1 Introduction

The Educational IoT (Internet of Things) has evolved into a heterogeneous ecosystem in the educational domain, consisting of smart classrooms, biometric attendance systems, surveillance cameras, interconnected laboratories, access-control systems and campus gateways. Recent studies in smart classrooms report that the technology is used inconsistently, evaluation is not consistent, and the design of the technology does not focus on security⁽¹⁾. A systematic analysis of IoT botnets reveals that DDoS, reconnaissance, spoofing, and brute-force and Mirai attacks are all constant threats and that attack behaviour and traffic diversity are significant factors affecting the likelihood of detection⁽²⁾. Although conventional classifier evaluations with CICIoT-2023 show high benchmark performance, this is based on the flow statistics of the dataset, which gives only limited insight into the transferability of the classifiers to sector-specific networks⁽³⁾. Laboratory traffic assumptions and dependence on class-distribution limit

the generalisation to operation, while two-stage clustering and deep classification enhance the separation of large attack categories⁽⁴⁾.

The dependency on source domain patterns hinders the use of attention-based processing of transformer-based multiclass detection, which comes with added computational overheads, yet offers good classification accuracy⁽⁵⁾. Although collaborative intrusion-detection frameworks decrease the centralised data exposure and enhance the privacy, the co-ordination of nodes with different heterogeneous nodes, communication overheads and trust management are yet to be resolved in open campus network⁽⁶⁾. Although hybrid autoencoder, LSTM and CNN architectures enhance the ability to recognize different kinds of attacks, model complexity, lack of role-level explanation, and lack of cross domain validation hinder the deployment of Edu-IoT in real-world scenarios⁽⁷⁾. Although distributed learning and privacy protection are possible with Federated Tab Transformer frameworks, the semantic mismatch between laboratory devices and educational device roles is not addressed by the heterogeneity of clients and the communication cost⁽⁸⁾. Although knowledge-distillation models alleviate the inference cost in resource-limited scenarios, compressed classifiers are still susceptible to the source domain bias when the behaviour of specific sectors is not included during representation learning⁽⁹⁾. Based on recent gap analyses, there is still no fully validated cross-environment behavioural transfer, and challenges to real-time operation, lightweight inference, adaptivity and explain ability/autonomous response remain.

Recent gap analyses of IoT intrusion-detection research identify real-time operation, lightweight inference, adaptability, explainability, and autonomous response as continuing deployment challenges⁽¹⁰⁾. A separate challenge arises during sector-specific deployment because high benchmark accuracy does not necessarily establish preservation of attack behaviour after device roles, communication structures, and operational semantics change. Educational IoT networks contain distinctive functional roles, including classroom surveillance units, biometric attendance systems, smart laboratory sensors, classroom equipment controllers, and campus gateways. Consequently, the central research gap concerns behavioural transfer rather than only attack-classification accuracy. CDBM-IDS addresses the gap through educational device-role abstraction, spectral behavioural representation, hierarchical deviation modelling, tri-context fusion, and sector-transfer consistency regularization.

2 Methodology

2.1 Overview of the Proposed Architecture

The proposed architecture identifies IoT cyberattacks by transferring behavioural knowledge learned from CICIoT-2023 to a derived educational-domain representation. As illustrated in Figure 1, the framework transforms raw IoT flow records into domain-aware behavioural embeddings. The pipeline combines statistical traffic information, educational device-role mapping, spectral behaviour, hierarchical deviation information, and sector-transfer consistency. EduTwin represents an educational interpretation of CICIoT-2023 traffic and does not contain independently captured traffic from a deployed educational network.

As illustrated in Figure 1, the process starts with ingesting the IoT flow records of the CICIoT-2023 dataset and then a data pre-processing phase, which carries out cleaning, normalization and feature derivation, to create standard features of flow-stats. The features are the result of two parallel processing paths. The former constructs an Educational IoT Digital Twin Layer (E-DTL), which maps the IoT devices in the laboratory to educational-sector functions and creates a graph of interactions which represents the pattern of campus IoT communication. This layer is a structural and behavioural replacement of an educational network.

At the same time, the normalised flow-stat attributes are passed through the Behaviour Spectrum Encoder (BSE) which encodes the dispersion and spectral properties of network traffic and permits to identify periodic, bursty and anomalous behaviours of communication related to IoT attacks. The output of E-DTL is then analysed by Hierarchical Deviance Encoder (HDE). It quantifies deviations between observed behaviours and learned centroids at the local (role) level and the global (network) level.

2.2 Dataset Construction and Preprocessing

2.2.1. Source Dataset and EduTwin Construction

CICIoT-2023, developed by the Canadian Institute for Cybersecurity, forms the source dataset for EduTwin generation. The benchmark contains traffic collected from a large-scale IoT testbed comprising 105 heterogeneous devices and 33 attack types grouped into seven major attack families, including DDoS, DoS, reconnaissance, web-based attacks, brute-force attacks, spoofing, and Mirai attacks⁽¹¹⁾. The distributed flow-level dataset contains 47 columns: 46 numerical predictive features and one class-label column. Accordingly, the present study treats the 46 numerical attributes as the original flow-level feature set, while the class label is used as the prediction target. EduTwin-specific role, zone, threat, and risk attributes are subsequently




Fig 1. Overall architecture of the proposed Cross-Domain Behavioural Mapping Intrusion Detection System (CDBM-IDS)

generated as domain metadata and are not counted as original CICIoT-2023 flow-level features. The official CICIoT-2023 dataset is available from the Canadian Institute for Cybersecurity at <https://www.unb.ca/cic/datasets/iotdataset-2023.html>.

The original CICIoT-2023 traffic is not treated as an unchanged educational-network benchmark. A derived dataset, termed EduTwin, is generated by transforming source traffic through educational device-role mapping and sector-oriented behavioural representation. Original IoT devices are mapped to functionally related educational roles, including classroom CCTV units, smart laboratory sensors, digital classroom equipment controllers, biometric attendance systems, and campus Wi-Fi access points or gateways. Source-destination information, protocol metadata, and communication frequency are subsequently used for construction of an educational interaction graph. Attack labels are preserved while laboratory-specific device semantics are replaced by educational-sector role semantics.

The final EduTwin evaluation dataset contains millions of flow records. For the binary intrusion-detection experiment reported in Section 4, malicious CICIoT-2023 attack labels are aggregated into a single Attack class, while normal traffic remains under the Benign class. The 70:30 stratified split produces a held-out test partition containing 58,140 flows: 29,915 benign flows and 28,225 attack flows. The test-set counts are directly consistent with the confusion matrix in Figure 4, where 29,574 benign flows and 27,812 attack flows are correctly classified, while 341 benign flows and 413 attack flows are misclassified.

2.2.2. Data Pre-Processing

Each CICIoT-2023 record contains numerical flow-level characteristics describing packet volume, packet timing, protocol usage, transmission rates, TCP flags, and statistical properties of the traffic. The 46 numerical flow-level features constitute the base input representation. The class-label column is excluded from the predictive feature matrix.

Data cleaning is performed before model training. Infinite values are converted to missing values, and invalid numerical entries are replaced with zero according to the implementation pipeline. Numerical features are subsequently standardised using z-score normalisation. Importantly, the mean and standard deviation are estimated only from the training partition and then applied unchanged to the test partition. Such training-only normalisation prevents information from the test data from influencing model preparation.

For feature j , z-score normalisation is defined as

$$x'_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j + \epsilon} \quad (1)$$

where x_{ij} denotes the value of feature j for flow i , μ_j and σ_j denote the mean and standard deviation of feature j calculated from the training partition, x'_{ij} denotes the corresponding standardised value, and ϵ denotes a small numerical constant used to avoid division by zero.

Following normalisation, the resulting flow-level representation is supplied to the Behaviour Spectrum Encoder and the statistical branch of the Tri-Context Fusion module. Educational role information generated by E-DTL is used separately for role-aware behavioural modelling and hierarchical deviation estimation.

After normalization, feature derivation for Behavioural Modelling (BM) is done. The features derived from BM are computed to support behavioural analysis. These involve aggregated activity measures, dispersion-based measures and role-independent flow measures that are common to all types of devices. Such derivations are important so that the model can concentrate on behavioural characteristics and not on device specific identities when the knowledge gained at CICIOT-2023 laboratory environment is applied to educational institutional networks.

The result of this process is a normalized feature matrix $X \in \mathbb{R}^{N \times d}$ with N being the number of network flows and d representing the number of normalized flow-level attributes. The rows of this matrix are associated with normalized feature vectors of a single IoT flow instance. This common feature matrix is the input of Behaviour Spectrum Encoder and Educational IoT Digital Twin Layer. To make spectral encoding, hierarchical deviation estimation and cross-domain behavioural mapping work on similar and stable representations as input, the preprocessing module offers feature scaling and consistency of representation of the flows.

2.3 Educational IoT Digital Twin Layer (E-DTL)

The E-DTL is a way to overcome the structure and contextual mismatch between the real world educational institutional networks and the real world and the laboratory based dataset. E-DTL reduces the possible errors that could occur because of transferring learned patterns to educational settings when using a controlled IoT testbed.

The E-DTL has two phases as illustrated in Figure 2, which include: device-role mapping and creation of the interaction graph. In the first stage, IoT devices in the dataset are remapped to the same kind of device role in the education sector based on functional similarity. The cameras become classroom CCTV cameras, the motion sensors become smart classroom sensors, the smart plugs become digital classroom equipment controllers, the doorbells become biometric attendance and IoT hubs become access points/gateways to campus Wi-Fi. This role-preserving remapping preserves behaviour patterns learnt in CICIOT-2023 flows semantically significant in the educational context.

Let D_{lab} be the set of IoT devices in the CICIOT-2023 environment and D_{edu} be the set of roles of educational-sector IoT devices. A mapping function is called deterministic if it is:

$$M : D_{lab} \rightarrow D_{edu} \quad (2)$$

This mapping allows to represent device identities through role based abstractions and the proposed method can focus on communication behaviour rather than hardware specific attributes.

The E-DTL in the second phase then builds an educational communication interaction graph with the mapped device roles and flow metadata derived out of the dataset. To build a graph, flow endpoints, protocol information and communication frequencies $G_{edu} = (V, E)$ are used, where $V = M(D_{lab})$ represents the roles of educational IoT devices and E represents the communication frequencies obtained from the source-destination flow relationships. Structural properties (connectivity patterns, communication density and intensity of interaction at a role-level) in the simulated educational network are captured in this interaction graph.

This structural backbone is the result of the E-DTL providing a context-sensitive way of behavioural modelling. Interaction graph is useful for subsequent modules such as Hierarchical Deviations Encoder to assess the deviations in not only flow level but also in device-role level and network level. The E-DTL is an important link between the laboratory IoT traffic and the requirements of educational institutions security, as shown in Figure 2, which enables many cross domain attacks to be identified.

2.4 Behaviour Spectrum Encoder (BSE)

The BSE learns the normal behavior patterns of devices and network traffic and constructs a spectrum of reference patterns of the normal behavior. It provides information on complex features of behaviour like frequency patterns, usage behaviour and




Fig 2. Educational IoT Digital Twin Layer

interaction styles which cannot be captured by time domain flow statistics. As can be seen in Figure 3, the BSE transforms the flow elements into a spectral form.

The BSE receives as input a normalized flow feature vector $X \in R^d$ where d represents the number of flow-level attributes that are extracted from the dataset as shown in Figure 3. The encoder represents a flow vector X as a behavioural signal, instead of a static snapshot of it. It employs Fast Fourier Transform (FFT) to get its frequency components. The Fourier transform of the vector X whose elements are $x_1, x_2, \dots, x_d$ is calculated as

$$S(k) = \text{FFT}(X) = \sum_{n=0}^d x_n e^{-j2\pi(k-1)(n-1)/d}, k = 1, 2, \dots, d \quad (3)$$

where

- x_n are the elements of the feature vector
- d is the total number of elements in X
- $S(k)$ is the frequency-domain output spectrum
- $j = \sqrt{-1}$
- k = index of the frequency component

The output spectrum S detects repetitions and periodicities of an attribute of the flow. These spectral signatures are very good in discriminating between bot controlled/organized attacks and normal educational traffic, which have high frequency domain regularities. The dispersion measure to measure the variability of the spectral components is:

$$D = \frac{\sqrt{1/d \sum_{k=1}^d (S(k) - \mu_S)^2}}{\mu_S + \epsilon} \quad (4)$$

Where $|S(k)| = \sqrt{(RP(S(k)))^2 + (IP(S(k)))^2}$ and $\mu_s = \frac{1}{d} \sum_{k=1}^d |S(k)|$

In the above equation for D , the numerator denotes standard deviation, μ_s is the mean of the spectral magnitudes and ϵ is a very small number to make the equation numerically stable. The dispersion D is the irregularity of the communication

behaviour of a device. The dispersion of S is measured so as to understand how uneven the frequency components are. The greater the dispersion, the more irregular is its behaviour and vice versa. Thus, the encoder allows the identification of benign patterns that are more likely to be stable and volatile, attack-driven traffic.




Fig 3. Workflow of BSE Layer

Then the BSE concatenates the original flow components X , the spectral components S and the dispersion measure D :

$$Z_{BSE} = [X, S, D] \quad (5)$$

These composite embedding retain both temporal and frequency domain information and can be used to provide a more expanded behavioral representation than traditional feature vectors. The obtained spectrum embedding Z_{BSE} is then sent to the Tri-Context Fusion Operator where it is fused with the hierarchical deviance information and the flow-stat features, as illustrated in Figure 3.

2.5 Hierarchical Deviance Encoder (HDE)

HDE detects deviant behaviours at multiple levels of the educational IoT digital twin structure. While the BSE focus on individual flows, it does consider the role of a device and its behaviour in comparison to the other devices in the network. The HDE does that by measuring the difference between the device-role norms and the network behaviour of the globe as illustrated in Figure 3.

Consider a behavioural fingerprint vector $F \in R^k$ corresponding to a particular flow f of its IoT device. It is a fingerprint based on flow-stat features obtained by the Educational IoT Digital Twin Layer and normalized together with the role-specific metadata. Every fingerprint F has a role r in the set of roles in graph G_{edu} of educational devices.

2.5.1. Role-Level Behavioural Centroid

For each educational device role r , a role-level behavioural centroid is computed using all fingerprints belonging to that role:

$$C_r = \frac{1}{N_r} \sum_{i=1}^{N_r} F_i \quad (6)$$

where:

- $C_r \in R^k$ is the centroid representing normal behaviour of role r ,
- N_r is the number of flows associated with role r ,
- F_i denotes the behavioural fingerprint of the i^{th} flow belonging to role r .

This centroid captures the typical communication behaviour expected from a specific educational IoT role, such as classroom CCTV units or smart laboratory sensors.

2.5.2. Global Network Behavioural Centroid

To model overall network behaviour, a global centroid is computed across all benign training flows:

$$\bar{C} = \frac{1}{N} \sum_{j=1}^N F_j \quad (7)$$

where:

- $\bar{C} \in R^k$ represents the global benign behavioural profile of the educational IoT network,
- N is the total number of benign flows used for centroid estimation,
- F_j denotes the behavioural fingerprint of the j^{th} benign flow.

The global centroid provides a reference point for detecting largescale anomalies affecting multiple device roles simultaneously.

2.5.3. Hierarchical Deviance Estimation

Using the computed centroids, the HDE evaluates two complementary deviation measures for each flow:

1. Local Role-Level Deviance

$$\Delta_1 = \|F - C_r\|_2 \quad (8)$$

This term quantifies how much the behaviour of a flow deviates from the expected behaviour of its assigned educational device role.

2. Global Network-Level Deviance

$$\Delta_2 = \|C_r - \bar{C}\|_2 \quad (9)$$

This term measures how far the role-level behaviour itself deviates from the overall network norm, enabling detection of coordinated or role-specific attack patterns.

Here, $\|\cdot\|_2$ denotes the Euclidean norm.

2.5.4. HDE Embedding Construction

The final hierarchical deviance embedding is formed by concatenating the two deviation components:

$$Z_{HDE} = [\Delta_1, \Delta_2] \quad (10)$$

where

- $Z_{HDE} \in R^2$ represents the hierarchical deviance embedding of a flow
- Δ_1 captures fine-grained role-specific anomalies
- Δ_2 captures broader network-level behavioural shifts.

2.6 Tri-Context Fusion Operator (TCF)

TCF is a single context based behavioural information that augment attack classification. The BSE provides information in the frequency domain, the HDE provides information in the network domain, and both types of information can be combined in an integrated and enriched embedding by the TCF as shown again in Figure. 3.

Let

- $X \in R^d$ denote the normalized flow-stat feature vector obtained after data pre-processing,
- $Z_{BSE} \in R^{d+d_f+1}$ denote the Behaviour Spectrum Encoder embedding, where d_f represents the number of frequency components derived from FFT,
- $Z_{HDE} \in R^2$ denote the Hierarchical Deviance Encoder embedding comprising local and global deviation measures.

Each of these vectors encodes distinct aspects of IoT communication behaviour: X reflects instantaneous flow statistics, Z_{BSE} captures periodicity and spectral dispersion and Z_{HDE} captures deviation from expected educational-role behaviour.

2.6.1. Geometric Fusion Mechanism

The TCF combines context-weighted geometric fusion rule in lieu of conventional concatenation or attention mechanisms to enhance the consistent behaviour signals across contexts and minimize the context-specific noise. The fused embedding is calculated as:

$$Z_{TCF} = (Z_{BSE}^{w_1} \odot Z_{HDE}^{w_2} \odot X^{w_3}) \frac{1}{w_1 + w_2 + w_3} \quad (11)$$

where

- Z_{TCF} is the fused tri-context embedding,
- $\odot$ denotes element-wise multiplication,
- $w_1, w_2, w_3 \geq 0$ are context-weight parameters controlling the relative influence of spectral, deviance and statistical information, respectively.

The exponent $\frac{1}{w_1 + w_2 + w_3}$ ensures scale normalization of the fused embedding and prevents dominance of any single context.

2.6.2. Context Weight Interpretation

The fusion weights provide interpretability and adaptability across deployment environments:

- w_1 emphasizes spectral behaviour, which is effective for detecting periodic attacks such as DDoS flooding and botnet propagation.
- w_2 emphasizes hierarchical deviations, supporting detection of role-specific and coordinated anomalies in educational IoT networks.
- w_3 preserves baseline statistical characteristics of network flows.

In practice, these weights are either empirically selected or optimized during training to maximize detection performance.

2.7 Sector-Transfer Consistency Learning (STC-Loss)

STC-Loss is designed to preserve representation consistency between the original CICIOT-2023 flow representation and the corresponding EduTwin educational-domain projection. E-DTL changes device semantics and introduces educational role information, which can alter the resulting representation even though the underlying traffic originates from CICIOT-2023. STC-Loss therefore constrains the source embedding and the derived educational projection during training. The resulting alignment should be interpreted as consistency across the original and transformed representations rather than validation across independently collected real-world networks.

Let

- $Z \in R^m$ denote the fused tri-context embedding obtained from the Tri-Context Fusion Operator for a given IoT flow in the CICIOT-2023 domain,
- $P(\cdot)$ denote the projection function induced by the Educational IoT Digital Twin Layer,
- $Z' = P(Z)$ denotes the corresponding embedding projected into the educational-sector representation space.

The objective of STC-Loss is to enforce representation invariance across the source domain (laboratory IoT dataset) and the target domain (educational IoT network) that preserves the discriminative power for attack classification.

2.7.1. STC-Loss Formulation

The proposed STC-Loss consists of two complementary components: a distance consistency term and a directional consistency term. The loss function is defined as:

$$\mathcal{L}_{STC} = \|Z - Z'\|_2^2 + \lambda(1 - \cos(Z, Z')) \quad (12)$$

where:

- $\|\cdot\|_2$ denotes the Euclidean norm,
- $\cos(Z, Z') = \frac{Z \cdot Z'}{\|Z\|_2 \|Z'\|_2}$ represents cosine similarity,
- $\lambda \geq 0$ is a regularization parameter controlling the relative importance of directional alignment.

The first term, $\|Z - Z'\|_2^2$ encourages that the embedding of the source-domain should not be substantially different from its projection in the educational-domain regarding the magnitude of the embeddings. This term can be added to make sure of the behaviour representation consistency in different levels of distance, so that the gap is not too large for the domain transfer.

The second term, $1 - \cos(Z, Z')$, is a penalization of angular mismatch between the two embeddings. This directional consistency constraint ensures that the relative orientation of embeddings in feature space is maintained, which is very important to preserve the class separability across the domains. In practical terms, this means that the model is not allowed to learn representations that look very different in the lab data but then rotate or collapse in an unpredictable fashion in the educational context.

2.7.2. Role of STC-Loss in Cross-Domain Learning

As shown in Figure 3, the STC-Loss is applied during training alongside the standard classification loss. By jointly optimizing both objectives, the model learns embeddings that are simultaneously

- **Discriminative**, enabling accurate classification of benign and attack traffic; and
- **Transfer-consistent**, ensuring that learned attack behaviours remain meaningful when mapped to educational institutional networks.

Since the dataset includes attacks created in a controlled setting that may not be representative of a real campus network, this dual-constrained methodology is essential. STC-Loss specifically addresses the gap through cross-domain behaviour matching, as opposed to using feature-similarity or retraining.

2.8 Classification Layer

The last element of the method is the classification layer. It classifies every IoT flow into its respective traffic category namely benign activity or a given kind of attack with the help of the single embedding generated by behavioural encoding, hierarchy deviance estimation, multi-context fusion and sector-transfer consistency regularization.

Let $Z_{TCF} \in R^m$ denote the fused embedding produced by the TCF operator. As illustrated in Figure 3, this embedding is provided as input to a lightweight dense neural classifier designed to balance discriminative capability and computational efficiency. The classifier applies a linear transformation followed by a softmax activation to compute class probabilities:

$$\hat{y} = \text{Softmax}(WZ_{TCF} + b) \quad (13)$$

where:

- $W \in R^{C \times m}$ is the weight matrix
- $b \in R^C$ is the bias vector
- C is the number of traffic classes, including benign traffic and multiple IoT attack categories defined in CICIoT-2023
- $\hat{y}$ represents the predicted probability distribution over classes.

The categorical cross-entropy loss is used to train the classifier and to promote accurate separation of benign and malicious traffic whilst maintaining inter-class separation between various types of attacks. This classification loss is trained simultaneously with Sector-Transfer Consistency Loss, to obtain both accuracy of the prediction and domain-consistency.

The use of a dense classifier instead of either a deeper or more complex decision module is not random. Much of the abstraction of behaviour and domain alignment is already done in the earlier stages of the method. This means that plain classifier will be adequate to attain high detection rate with low inference overhead. This design will come in handy when applied in real networks in the education sector, where real time detection and limited computing resources are significant.

2.9 Experimental Setup

The experimental configuration is summarised in Table 1. CICIoT-2023 provides 46 numerical flow-level predictive features together with one class-label column. For the reported binary experiment, all malicious labels are grouped into the Attack class and normal traffic is retained as the Benign class. A 70:30 stratified train-test split is applied using random seed 42. Numerical features are standardised by fitting z-score parameters exclusively on the training partition. CDBM-IDS is trained for 25 epochs using Adam with a learning rate of 0.001 and a batch size of 2048. The fusion and backbone dimensions are 128, while the role embedding dimension is 16. A dropout rate of 0.10 is applied in the backbone. STC-Loss uses an overall weight of 0.50 and a cosine-consistency coefficient of 0.50. All three TCF context weights are fixed at 1.0. Experiments are executed using PyTorch with CUDA acceleration when an NVIDIA GPU is available.

Table 1. Experimental setup and implementation configuration for CICIoT-2023 and EduTwin evaluation

Parameter	Configuration
Source Dataset	CICIoT-2023
Derived Dataset	EduTwin
Classification	Binary: Benign vs Attack
Data Split	70:30 stratified train-test split
Normalization	Z-score standardization
Optimizer	Adam
Learning Rate	0.001
Epochs	25
Batch Size	2048
Fusion Dimension	128
Role Embedding Dimension	16
Dropout Rate	0.10
STC-Loss Weight	0.50
STC Cosine Coefficient, (α)	0.50
TCF Weights	$w_b = w_h = w_x = 1.0$
Random Seed	42
Evaluation Metrics	Accuracy, Precision, Recall, F1-score, AUC-ROC
Implementation	PyTorch with CUDA-enabled GPU

3 Results and Discussion

This section discusses the experimental outcomes achieved using the CICIoT-2023 dataset and analyzes the performance of the proposed CDBM-IDS method. The evaluation is based on standard classification metrics, along with extra robustness metrics, statistical significance testing and comparative analysis with popular baseline models.

3.1 Confusion Matrix Analysis

The confusion matrix of the intrusion detection model developed based on the EduTwin is shown in Figure 4, and it shows the model has a high capability to differentiate the benign traffic and the attack traffic. The model accurately detected 29,574 benign flows in the matrix and falsely labeled 341 benign flows as an attack in it- a very low rate of false-positive which is crucial in reducing false alarms in a network environment based on education. In the case of attack traffic, 27812 attack flows were correctly identified and only 413 attack examples were incorrectly identified as benign. This even distribution of true positive and true negative ratio signifies that the proposed model is reliable and strong system in attack detection and can face the impact of false alarms and is useful in protecting the heterogeneous educational IoT infrastructures.

3.2 Comparative Analysis with Baseline Models

Table 2 presents a controlled comparison between CDBM-IDS and conventional deep-learning architectures executed under the same EduTwin experimental protocol. DNN obtains 95.02% accuracy, while recurrent architectures such as LSTM, BiLSTM,




Fig 4. Confusion Matrix of the Proposed Model

and GRU obtain accuracies between 96.47% and 96.83%. CNN-LSTM and CNN-BiLSTM improve accuracy to 97.11% and 97.36%, respectively. Transformer and TCN obtain 97.54% and 97.68% accuracy, while GNN produces the strongest reproduced baseline performance with 97.81% accuracy, 97.66% precision, 97.49% recall, 97.57% F1-score, and 98.84% AUC-ROC.

CDBM-IDS obtains 98.72% accuracy, 98.64% precision, 98.51% recall, 98.57% F1-score, and 99.21% AUC-ROC. Relative to GNN, CDBM-IDS improves accuracy by 0.91 percentage points, precision by 0.98 percentage points, recall by 1.02 percentage points, F1-score by 1.00 percentage point, and AUC-ROC by 0.37 percentage points. The controlled comparison indicates additional discriminative benefit from behaviour-spectrum representation, hierarchical deviation modelling, geometric context fusion, and sector-transfer consistency learning under the EduTwin representation. The comparison does not imply numerical superiority over separately published CICIOT-2023 studies because external studies use different traffic representations and experimental protocols.

Table 2. Comparative Performance Analysis

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
DNN (Fully Connected)	95.02	94.81	94.36	94.58	96.12
CNN-1D	96.21	96.04	95.83	95.93	97.18
CNN-2D	96.48	96.31	96.02	96.16	97.44
LSTM	96.47	96.31	96.02	96.16	97.46
BiLSTM	96.83	96.68	96.41	96.54	97.88
GRU	96.59	96.43	96.19	96.31	97.63
CNN-LSTM Hybrid	97.11	96.94	96.78	96.86	98.07
CNN-BiLSTM Hybrid	97.36	97.21	97.02	97.11	98.34
Autoencoder + DNN	97.02	96.88	96.59	96.73	98.01
Transformer (Flow-based)	97.54	97.39	97.18	97.28	98.62
Temporal Convolutional Network (TCN)	97.68	97.53	97.34	97.43	98.71
Graph Neural Network (GNN)	97.81	97.66	97.49	97.57	98.84
Proposed CDBM-IDS	98.72	98.64	98.51	98.57	99.21

3.3 Statistical Significance Analysis

A statistical significance test that can analyze whether the performance improvement of the proposed CDBM-IDS model compared to the best baseline (Autoencoder+DNN) is real or a result of random fluctuation is provided in Table 3. The paired t-test resulted in t values of 6.37 and 5.94 for Accuracy and F1 score respectively with p values of less than 0.001 for both, which rejected the null hypothesis with high level of confidence. Accuracy and F1-score had Wilcoxon signed-rank test of Z-scores of -4.89 and -4.61 respectively with p-values less than 0.001 which confirms that the improvements do not disappear in non-parametric test environment. In general, these findings have shown that the CDBM-IDS model benefits statistically significant improvements at the 95% confidence, which proves to be better than the baseline in terms of detection accuracy and even balanced classification detection.

Table 3. Statistical significance analysis comparing the proposed CDBM-IDS

Metric	Paired t-Test		Wilcoxon Signed-Rank Test	
	t-value	p-value	Z-score	p-value
Accuracy	6.37	< 0.001	-4.89	< 0.001
F1-Score	5.94	< 0.001	-4.61	< 0.001

3.4 Ablation Study

From the ablation study in Table 4, it is observed that all the proposed components play a crucial role in the detection accuracy and robustness. Eradication of BSE also comes at a cost in terms of performance, since it wipes out the frequency-domain behavioural clues that identify periodic bot-net and DDoS activity. The decreased recall due to HDE dropping is because role-aware deviation modelling is critical to the detection of small anomalies in device-role clusters. With TCF replaced by some simple concatenation, false positives are more eye-catching and geometric fusion is important to reduce context-dependent noise and to amplify consistent attack evidence in a cross-context fashion. The maximum decrease is seen in the scenario where the STC -Loss is removed, indicating that the learning of sector-transfer consistency is crucial to make the CICIoT 2023 representations to the educational IoT digital twin and improve the cross-domain generalisation.

3.5 Discussion

The proposed work captures frequency-based regularities and role-conscious behavioral anomalies which are identified to be the most critical ones in detecting the presence of coordinated IoT attacks, while most deep learning based models are based on time-based regularities. The introduction of STC Loss improves the generalisation as it aligns with laboratory information and

Table 4. Ablation Study of CDBM-IDS Components on CICIoT-2023

Variant	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)	FPR (%)	MCC
Full CDBM-IDS (BSE + HDE + TCF + STC-Loss)	98.72	98.64	98.51	98.57	99.21	1.14	0.973
w/o BSE (remove spectral embedding)	97.86	97.71	97.44	97.57	98.63	1.62	0.956
w/o HDE (remove hierarchical deviation)	97.63	97.52	97.19	97.35	98.51	1.74	0.951
w/o TCF (replace fusion with simple concatenation)	97.41	97.29	96.98	97.13	98.32	1.89	0.946
w/o STC-Loss (no sector-transfer consistency)	97.12	96.98	96.71	96.84	98.04	2.11	0.940

the real world educational settings. This allows the deployment to be performed reliably, without the need for labelled attack data on real campus networks.

Although Transformers and GNNs are popular on network security but they do not perform as well as proposed CDBM-IDS on the modified CICIoT-2023 dataset. Their strength is limited by the characteristics of the data, assumptions made in performing the model analysis and limitations in their application.

Self-attention is used to learn long-range dependencies in sequences using transformers. CICIoT-2023 does not provide real sequences, but summary flow-based records of packet activity. The expectation for a high-order assumption is then low and redundant or weak features, that can only give marginal improvement, are emphasized, at the price of greater complexity and training instability.

CDBM-IDS does not assume strict time continuity and fixed graphs. It concentrates on behavior-based abstractions that are not sensitive to ordering and change of topology. BSE can record periodic and burst patterns that are not recorded by Transformers at the flow level. The HDE detects graphless role anomalies and graphless network anomalies. The TCF operator also eliminates noisy signals which are context-dependent and cause bias in attention or message-passing models.

Transformers or GNNs have high latency for inference and high memory requirements and cannot be used in real-time monitoring in resource-limited educational settings. The lightweight classification and sector-transfer consistency learning of CDBM-IDS provides a superior tradeoff between accuracy, generalization and operational feasibility.

4 Conclusion

In the field of IoT, a new cross domain intrusion detection system is presented in this study which is being implemented to an education IoT system. It brings together generic IoT attack intelligence, as well as educational context-specific awareness. The authors created an educational based benchmark, called EduTwin, based on the CICIoT-2023 dataset. What is seen as a major gap is that most of the detection models are not knowledgeable of a domain and cannot be easily interpreted in the actual world of education. The CDBM-IDS method is a fusion of two techniques: behavioural encoding and hierarchical deviance modelling and tri -context fusion. It will be able to identify good and bad traffic, accurately and still function optimally in other sectors. The model was extensively tested and compared to numerous classical machine learning, deep learning and state-of-the-art architectures such as CNN -BiLSTM, Transformer models, TCNs and GNNs. CDBM-IDS obtained 98.72% accuracy, 98.57% F1-score and 99.21% AUC-ROC. In complex and unbalanced attack situations these values show good detection. Improvements proved to be significant (using statistical tests such as paired t -tests and Wilcoxon signed -rank tests) with a confidence level of 95% which is better than the best baseline. The research also included raw numbers as well as strength and practicality. This is especially useful when it involves the education of resources of the greatest significance, like data centers, administration servers and surveillance systems. While it is important to achieve high detection rates it is also important to minimise the number of false alarms and to safeguard critical infrastructure. The work is in progress of flow-level and offline analysis. The future research will be focused in real-time implementation, risk-sensitive thresholds and sparse edge computing solutions to resource-constrained educational IoT devices.

5 Acknowledgements

The author would like to express sincere gratitude to St. Joseph's College Research Grant - 2025 (SJCRG-2025) for providing the necessary funding that made this research possible and DST-FIST, Government of India for funding towards Infrastructure facilities at St. Joseph's College (Autonomous), Tiruchirappalli-620002.

Author details: Basil Raja P is a Research Scholar; Dr. L. Arockiam is a an Associate Professor; Dr. A. Aloysius is an Assistant Professor.

References

- 1) Figueroa E, Batista E, Palau R, Unciti O, Ferre M, Martinez-Balleste A. The use of artificial intelligence techniques in smart classrooms is in its infancy. *IEEE Access*. 2024;12:125179–125193. <https://doi.org/10.1109/ACCESS.2024.3454372>.
- 2) Gelgi M, Guan Y, Arunachala S, Rao MSS, Dragoni N. Systematic literature review of IoT botnet DDoS attacks and evaluation of detection techniques. *Sensors*. 2024;24(11):3571. <https://doi.org/10.3390/s24113571>.
- 3) Kumar AG, Rastogi A, Ranga V. Evaluation of different machine learning classifiers on new IoT dataset CICIoT2023. *Proceedings of the 2024 International Conference on Intelligent Systems for Cybersecurity*. 2024;p. 625–630. <https://doi.org/10.1109/ISCS61804.2024.10581375>.
- 4) Ghani HQ, Al-Yaseen WL. Two-step data clustering for improved intrusion detection system using CICIoT2023 dataset. *e-Prime—Advances in Electrical Engineering, Electronics and Energy*. 2024;9:100673. <https://doi.org/10.1016/j.prime.2024.100673>.
- 5) Tseng SM, Wang YQ, Wang YC. Multi-class intrusion detection based on Transformer for IoT networks using CIC-IoT-2023 dataset. *Future Internet*. 2024;16(8):284. <https://doi.org/10.3390/fi16080284>.
- 6) Wardana AA, Kołaczek G, Sukarno P. Lightweight, trust-managing, and privacy-preserving collaborative intrusion detection for Internet of Things. *Applied Sciences*. 2024;14(10):4109. <https://doi.org/10.3390/app14104109>.
- 7) Susilo B, Muis A, Sari RE. Intelligent intrusion detection system against various attacks based on a hybrid deep learning algorithm. *Sensors*. 2025;25(2):580. <https://doi.org/10.3390/s25020580>.
- 8) Elaziz MA, Fares IA, Dahou A, Shrahili M. Federated learning framework for IoT intrusion detection using TabTransformer and nature-inspired hyperparameter optimization. *Frontiers in Big Data*. 2025;8:1526480. <https://doi.org/10.3389/fdata.2025.1526480>.
- 9) Wisanwanichthan T, Thammawichai M. A lightweight intrusion detection system for IoT and UAV using deep neural networks with knowledge distillation. *Computers*. 2025;14(7):291. <https://doi.org/10.3390/computers14070291>.
- 10) Sallam S, Barachi ME, Li N. Intrusion detection on the Internet of Things: A comprehensive review and gap analysis toward real-time, lightweight, adaptive, and autonomous security. *IoT*. 2026;7(1):16. <https://doi.org/10.3390/iot7010016>.
- 11) Neto ECP, Dadkhah S, Ferreira R, Zohourian A, Lu R, Ghorbani AA. CICIoT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment. *Sensors*. 2023;23(13):5941. <https://doi.org/10.3390/s23135941>.