Online-Journals.org

iJEP International Journal of Engineering Pedagogy

iJEP | eISSN: 2192-4880 | Vol. 16 No. 6 (2026) | OPEN ACCESS

https://doi.org/10.3991/ijep.v16i6.62878

PAPER

Privacy-Preserving Edge AI with Graph Transformer for Secure Real-Time Smart Engineering Education in Dense Urban Areas

Haider TH. Salim ALRikabi[1] (✉), Hussain Ali Mutar[1], Hind Sabeeh Raheem[2], Huda Abbas Kanber[2], Sanaa Aliwi Abdel Sada Habib[2], Ibtihal Razaq Niama ALRubeei[1]

[1]Wasit University, Wasit, Iraq

[2]University of Baghdad, Baghdad, Iraq

hdhiyab@uowasit.edu.iq

ABSTRACT

Engineering education is becoming more advanced in rapidly expanding urban centers with smart connected classrooms, laboratories, wireless networks, indoor environment sensors, and digital learning platform technologies. This establishes three major challenges and requires real-time educational analytics, individual campus data protection, and a robust, deployable system capable of withstanding cyber risks, heavy congestion & heterogeneous edge hardware. However, the majority of existing literature on edge analytics and federated learning or graph-based models, or security monitoring , addresses them individually with limited attention to unifying real-time laurence of secure architectures supporting multi-buildings on engineering campuses. This paper proposes a privacy-preserving edge-graph transformer (PP-EGT) for secure real-time smart engineering education in dense urban environments. The model integrates local anonymization, temporal encoding, dynamic campus graph construction, graph-transformer reasoning, and federated secure aggregation. It jointly predicts a five-class zone state, a continuous Learning Readiness Index (LRI), and a cyber-physical incident probability using multimodal edge features from environmental, network, pedagogical, and security domains. A literature-grounded simulation benchmark representing 48 zones across six campus buildings was constructed over 12 weeks with 5 min sampling. Across five random seeds, PP-EGT achieved 91.89% macro-F1, 96.45% AUROC for incident detection, and 3.43 LRI MAE. Compared with the strongest baseline, ST-Transformer, PP-EGT improved macro-F1 by 1.45 points, reduced p95 latency from 49.59 ms to 44.02 ms, and decreased privacy-leakage AUC from 69.19% to 53.83%. Its novelty lies in jointly optimizing learning analytics, cyber-physical security, privacy, and edge deployment constraints.

KEYWORDS

smart engineering education, edge AI, graph transformer, federated learning, privacy-preserving analytics, dense urban campus

ALRikabi, H. T. S., Mutar, H. A., Raheem, H. S., Kanber, H. A., Habib, S. A. A. S., ALRubeei, I. R. N. (2026). Privacy-Preserving Edge AI with Graph Transformer for Secure Real-Time Smart Engineering Education in Dense Urban Areas. *International Journal of Engineering Pedagogy (iJEP)*, 16(6), pp. 143–163. https://doi.org/10.3991/ijep.v16i6.62878

Article submitted 2026-06-09. Revision uploaded 2026-08-03. Final acceptance 2026-08-04.

© 2026 by the authors of this article. Published under CC-BY.

1 INTRODUCTION

Engineering campuses located in a densely populated urban area are typified not as isolated classrooms but as being part of an evolving cyber-physical ecosystem for learning. Using video-enabled laboratories and the latest environmental sensors, wireless access points, access-control devices, data-real-time learning platforms, and digital twin dashboards provide ongoing access to data that measure indicators of teaching quality and student engagement, safety, and infrastructure health [1–8].

As urban density increases due to population growth in Asia and Africa, many students attending these engineering institutions will become more engaged with educational technology as they complete their engineering degrees through laboratory-based teaching and project-based learning; therefore, both on-campus and online learning will require immediate support rather than relying on cloud-only analytics. Reports by UNESCO and NIST stress the need for human-centric, accountable AI in education and the importance of building trust in the privacy management of individuals and continuous governance of AI-enabled educational systems [9–15].

Recent research has investigated new models for electrical engineering concepts, such as edge learning analytics, federated smart campuses, graph neural networks (GNN) and graph transformers, to mention a few, although there are still several challenges that researchers must overcome before implementing these practices at scale. For example, current implementations of GNNs or edge learning Analytics do not jointly consider the development of a model for pedagogical states together with network degradation and cyber-physical risks. In addition, the issues of privacy and confidentiality of both students and employees are seldom examined in a controlled manner as per standards set forth in other research studies. Finally, the existing computer-based simulations of educational environments ignore many of the critical relationships among the labs, lab schedules, wireless contention, and instructor movement [16–19].

This paper proposes a new privacy-preserving Edge Graph-Transformer architecture to facilitate real-time compliance with secure Smart Campus environments while allowing the management of latency, energy consumption, and privacy constraints. The key contributions of this work include a dynamic campus graph formulation, federated privacy-preserving graph transformer models, reproducible simulation protocols, evaluation of the architecture against the following criteria accuracy, latency, energy consumption, robustness, privacy, and statistical testing. The Challenge-to-design mapping in the proposed study is shown in Table 1.

Table 1. Challenge-to-design mapping in the proposed study

Challenge	Manifestation in Campus Operation	Design Response
Dense urban congestion	higher latency, packet loss, AP overload	edge inference + dynamic A_wifi,t
Privacy-sensitive educational data	surveillance and data-retention risk	local anonymization + DP-SGD + secure aggregation
Relational campus effects	single-room models miss cross-zone spillovers	dynamic campus graph + graph transformer
Operational decision pressure	need for actionable alerts under ms constraints	multi-task heads + latency-constrained optimization

2 RELATED WORK

The focus of the scoping review includes education technology, Edge Artificial Intelligence (AI), Federated Learning, Privacy Engineering, and Graph Learning; therefore, the scoping focus was on cross-disciplinary synthesis, not a formal Registered Systematic Review. The academic sources indexed were combined with standards, guidance documents, etc., representing both methodology and governance perspectives for this area of study. The targeted scoping-review protocol used to synthesize recent literature is shown in Table 2.

Edge Analytics (EA) in Education represents the first cluster. In work conducted by Amo-Filva et al. on Privacy in Learning Analytics within Fog and Edge Computing, it was noted that Privacy by Design still lacks sufficient development, regardless of the advantages of Decentralized Processing for Education [20]. Valdes-Ramirez et al. showed how local classroom processing has a great reduction of bandwidth, yet still maintains anonymity within in-person instruction [21, 22]. Additionally, the findings of Kong et al. suggested that Spatiotemporal Graph Reasoning in Psychological Analytics of Classroom Behavior was improved using EA's Graph Model; they still concentrated on Recognizing Behaviors and Evaluating the Quality of Teaching and did not provide any evidence towards the secure multi-building coordination needed for improved outcomes [23].

The second cluster is represented by Privacy-Preserving Smart-Campus (S-C) & Federated Systems and includes work from Geng et al., testing the model of FedCampus (S-C – Federated Learning Application + Analytics using Privacy Protected Smartwatch Data) [24]. Mahmud et al. and Aminifar et al. highlighted that the use of federated training methods improves privacy in cyber-physical systems (CPS) and resource-constrained Internet of Things (IoT), at the same time, keeping the same level of detection performance [25, 26]. Although these studies present scalable privacy mechanisms, they do not align directly with the engineering-education decision loop model [27].

Cluster three consists of relational and graph-based educational intelligence. Pu et al. provided (survey) evidence on the benefits of using graph neural networks within intelligent tutoring systems. They presented that prerequisite structures, interaction networks, and resource dependencies all provide pedagogical advantages of utilizing graph models [28]. The research on GRapH-Transformers has also gained popularity, based on surveys, which suggest that they perform well on node-, edge- and graph level tasks with an increasing focus on scalability and robustness [29]. STGformer highlights the fact that Spatiotemporal Graph-Transformer Design Models provide strong predictive performance and efficient inference, while dealing with large-scale structured data [30].

The Fourth and Final Cluster relates to Privacy-Aware Urban AI and Governance. Ivanovska et al. established a Privacy-By-Design (PbD) AIOT Vision for Intelligent Urban Settings, which focuses on Embedded, Privacy-Protected Acquisition versus Post-Hoc redaction [31]. This approach is particularly pertinent to dense urban campuses, due to privacy concerns of sensing (e.g. video streams and access control streams). On a governance level, both the United Nations Educational, Scientific and Cultural Organization (UNESCO) and the National Institute of Standards and Technology (NIST) maintain that educational and urban AI systems need to be held accountable, require privacy management, and have human-In-the-loop safeguards [9, 32, 33].

The synthesis of the examinations presents a clear gap that remains. The research arguments are made for all four clusters: Edge Analytics, Federated Privacy, Graph-Based Educational Reasoning, and Urban Privacy Engineering, but there is no unified model to address real-time pedagogy, security risk, and edge deployment constraints for dense urban engineering education. In order to address this gap, it is

proposed that the PPIE be employed to create an end-to-end, measurable architecture that combines the strands.

Table 2. Targeted scoping-review protocol used to synthesize recent literature

Databases/Sources	Years	Search String (Condensed)	Inclusion Criteria	Exclusion Criteria	Yield
Scopus, Web of Science, ACM DL, arXiv, NIST, UNESCO, IETF	2020–2026	("smart campus" OR "smart classroom" OR "learning analytics") AND (edge OR federated OR privacy) AND (graph OR transformer OR GNN)	English; smart-campus or smart-education; privacy/security relevance; real-time or near-real-time analytics	non-education consumer apps; purely cloud-only dashboards without analytics; no methodological details	238 identified/161 after dedup/72 full-text screened/24 used in synthesis

The Taxonomy map used to organize the related-work discussion shown in Figure 1.

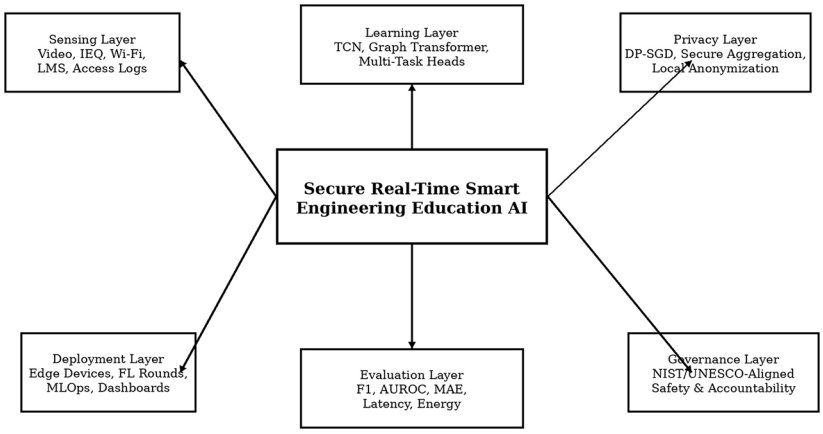


Fig. 1. Taxonomy map used to organize the related-work discussion

A comparative summary of representative prior work and the remaining gap showing in Table 3.

Table 3. Comparative summary of representative prior work and the remaining gap

Study	Primary Setting	Execution Mode	Strength	Gap
Amo-Filva et al. (2024) [20]	learning analytics privacy review	fog/edge	policy and architectural review	no graph modeling; no deployment metrics
Valdes-Ramirez et al. (2025) [21]	classroom behavior + IEQ	edge	local video processing, strong bandwidth reduction	single-task analytics; limited cross-building coordination
Geng et al. (2024) [24]	smart campus FL app	federated mobile	DP-enabled smart-campus analytics	not graph-aware; limited campus control-loop focus
Pu et al. (2024) [34]	ITS survey	graph neural networks	shows educational value of graph structure	survey, not real-time edge deployment
Kong et al. (2025) [23]	smart classroom GNN	spatiotemporal GNN	strong behavior modeling	privacy and FL support remain limited
Mahmud et al. (2024) [25]	IoT intrusion detection	federated privacy	security focus and decentralized training	not education-specific; no pedagogical outputs
PP-EGT (this work)	urban smart engineering education	edge + federated graph transformer	joint pedagogy-security inference with privacy and latency constraints	simulation-based only; no live pilot

3 PROBLEM FORMULATION

We model a smart engineering campus at time t as a dynamic graph $G_t = (V, E, A_t)$. The node set V comprises instructional zones—laboratories, lecture rooms, makerspaces, shared hubs—while edges E represent permissible inter-zone relationships and A_t is a time-varying weighted adjacency matrix. Each node $i \in V$ hosts an edge-enabled sensing and inference unit. The learning task is to infer, for each zone, a local pedagogical-operational state and an associated risk level from multimodal observations, under explicit latency, energy, and privacy constraints.

The observation at node i and time t is a concatenated feature vector

$$x_{i,t} = [e_{i,t} \| n_{i,t} \| p_{i,t} \| s_{i,t}] \in \mathbb{R}^F,$$

Where:

- $e_{i,t}$ contains environmental features (CO_2, $PM_{2.5}$, temperature, humidity, noise, and illuminance),
- $n_{i,t}$ contains network features (RSSI, access-point load, latency, jitter, packet loss),
- $p_{i,t}$ contains pedagogical features (attendance delta, LMS click-through rate, quiz response rate, device usage, participation proxies),
- $s_{i,t}$ contains security features (failed logins, abnormal port scans, MAC churn, and anomalous traffic ratio).

The primary label $y_{i,t}$ is a five-class zone state: *normal*, *under-engaged*, *environmentally stressed*, *network congested*, or *security critical*. Two auxiliary targets are also predicted: a Learning Readiness Index $r_{i,t} \in [0,100]$ and a binary indicator $z_{i,t} \in [0,1]$ for cyber-physical incident likelihood.

Dynamic graph construction. The time-varying adjacency matrix is a learned softmax combination of three components:

$$A_t = \text{softmax}(\alpha A_{\text{phys}} + \beta A_{\text{wifi},t} + \gamma A_{\text{ped},t}),$$

where A_{phys} encodes spatial proximity, $A_{\text{wifi},t}$ reflects current wireless contention and signal overlap, and $A_{\text{ped},t}$ captures pedagogical affinities such as shared courses, synchronized labs, and instructor circulation patterns. The scalars α, β, γ are non-negative mixing coefficients learned during training. This formulation allows the model to leverage both persistent infrastructure relationships and transient inter-zone dynamics.

Model architecture. Given a temporal window of length w, a local temporal encoder first processes each node's history:

$$h_{i,t}^{(0)} = \text{TCN}\left(x_{i,t-w+1:t}\right).$$

The resulting node embeddings are refined through a graph-transformer stack, where each layer applies multi-head attention over the dynamic graph, followed by residual connection, layer normalization, and a position-wise feed-forward block:

$$H_t^{(l+1)} = \text{LN}(H_t^{(l)} + \text{MHA}(A_t, H_t^{(l)})), H_t^{(l+1)} = \text{LN}(H_t^{(l+1)} + \text{FFN}(H_t^{(l+1)})).$$

From the final embeddings, task-specific heads produce the zone state $\hat{y}_{i,t}$, the readiness index $\hat{r}_{i,t}$, and the incident likelihood $\hat{z}_{i,t}$.

Federated learning objective. Training is federated across K edge clients, each hosting one or more campus zones. The local loss on client k is a weighted multi-task objective:

$$\mathcal{L}_k(\theta) = \lambda_1 \, \mathrm{CE}(\hat{y}, y) + \lambda_2 \, \mathrm{MAE}(\hat{r}, r) + \lambda_3 \, \mathrm{BCE}(\hat{z}, z),$$

where CE is cross-entropy, MAE the mean absolute error, BCE binary cross-entropy, and $\lambda_j \geq 0$ user-defined weights. The global optimization problem is

$$\min_{\theta} \sum_{k=1}^{K} w_k \, \mathcal{L}_k(\theta),$$

subject to practical deployment constraints:

- p_{95} inference latency ≤ 60 ms,
- per-inference edge energy $E_{\mathrm{inf}} \leq 250$ mJ,
- differential-privacy budget $\epsilon \leq 3.0$,
- packet-delivery ratio ≥ 99%.

These constraints convert the problem from a purely academic prediction exercise into a deployment-conscious engineering task, systematically surfacing the trade-offs often left implicit in educational AI research.

Assumptions. The formulation rests on four operational premises. (i) All raw personally identifiable data remain on-site; only anonymized features or privacy-protected model updates leave the edge node. (ii) Building clocks are synchronized sufficiently to align 5-minute observation windows. (iii) Graph edges are sparse and locally meaningful rather than fully connected. (iv) Ground-truth labels are generated by a rule-based simulator with stochastic perturbations that mimic realistic operational uncertainty. The present evaluation is therefore limited to simulation and does not claim human-subject validation. The dynamic campus graph used in the problem formulation is shown in Figure 2.

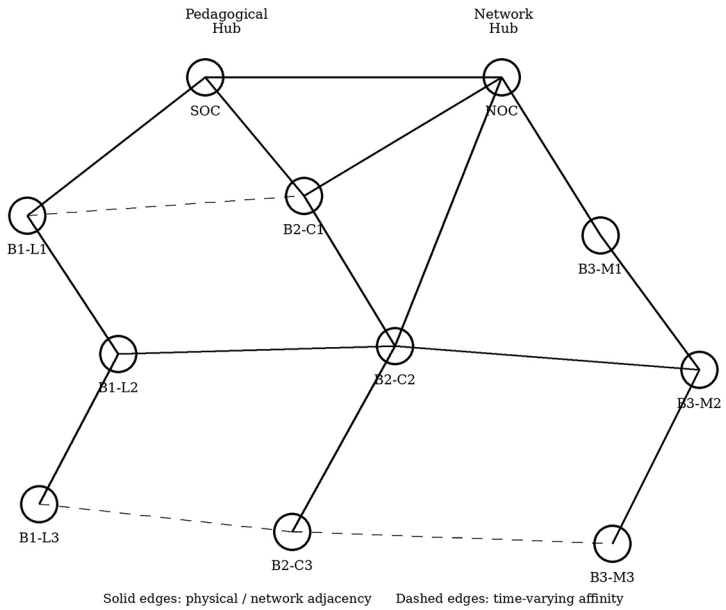


Fig. 2. Dynamic campus graph used in the problem formulation

The Core notation used in the problem formulation illustrated in Table 4.

Table 4. Core notation used in the problem formulation

Symbol	Meaning
$G_t = (V, E, A_t)$	Dynamic campus graph with nodes, edges, and weighted adjacency
$x_{i,t}$	Multimodal feature vector at zone i, time t
A_t	Dynamic adjacency from physical, wireless, and pedagogical affinities
$y_{i,t}$	Five-class zone state label
$r_{i,t}$	Learning Readiness Index (0–100)
$z_{i,t}$	Cyber-physical incident likelihood (0/1)
θ_k, θ	Local and global model parameters
w_k	Normalised client contribution to federated objective
ϵ	Differential-privacy budget
w	Temporal window length (number of 5-min slots)

4 PROPOSED METHOD/SYSTEM

The proposed PP-EGT system is organized as a four-stage pipeline. Stage 1 performs local sensing and anonymization. Stage 2 converts multimodal streams into compact temporal embeddings. Stage 3 performs dynamic graph-transformer reasoning over inter-zone dependencies. Stage 4 executes federated privacy-preserving optimization and real-time inference. This organization keeps personally identifiable raw data at the source while still allowing campus-wide learning across buildings.

How anonymity is achieved is done through the anonymous anonymization of private information. All video streams are reduced to count, trajectory, posture summaries, and occupancy descriptors. All access events are tokenized. All identifiers are replaced with salted hashes in a manner consistent with privacy-aware edge classroom systems and privacy-by-design urban AIoT work, as shown in Figure 3. Rather than simply complying with laws, a second objective is to minimize the architectural impact on users' privacy.

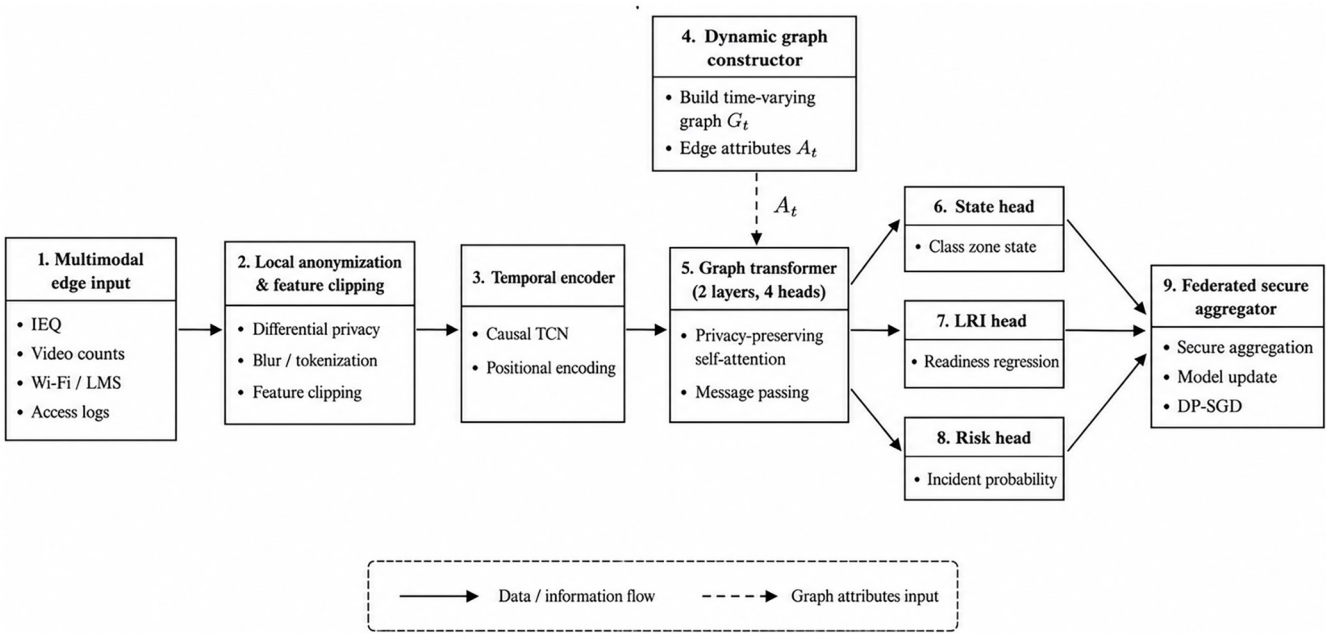


Fig. 3. PP-EGT architecture for privacy-preserving real-time smart engineering education

4.1 Components: temporal encoding, graph reasoning, privacy, and heads

The temporal encoder uses a lightweight causal temporal convolution network (TCN). This type of network allows for parallel computation and provides stable fixed-latency inference, which is advantageous for deployment to edge devices, compared to using recurrent baselines. As such, the role of the TCN is to compress information from the last few minutes of multi-modal context into a local hidden state while maintaining the ability to model short-term temporal events such as sudden changes in air quality, surges in class start times, and burst traffic of unanticipated events. For balancing the need for representational capacity and the limitations of edge memory, the output dimension of hidden state in the TCN is fixed at 64.

The graph transformer utilizes the dynamic nature of the campus graph through reasoning. The adjacency matrix used to compute attention is also weighted dynamically based on space, wireless, and pedagogical relationships; as a result, the structure of a GCN or GAT will not give a truly accurate depiction of how these relationships affect events on campus. Due to the nature of urban engineering campuses with non-Euclidean relationships, these types of activities/patterns do exist. For example, two labs on different floors may be utilizing the same overloaded access point to send data simultaneously, while two classrooms may load similar LMSs during a quiz, and critical points during emergency rerouting may appear to be the same throughout the entire building. Therefore, the graph transformer serves as not only an accuracy tool but also as a means of relational reasoning as a function of the design of the problem [35].

On top of the representation shared by the task heads, three heads are connected. The state head generates a five-class operational/pedagogical state. The second head, LRI, generates a bounded regression value that can be used to summarize short-term readiness to perform effective engineering instruction. The third head, which assesses risk, generates a probability of the occurrence of a cyber-physical event, allowing the model to function as both an educational and security-aware monitor. The use of multi-task learning facilitates the representation, being able to represent the combined structure of comfort, connectivity, engagement, and risk, as opposed to optimizing each component in isolation.

The ability to protect privacy is embedded in the training process rather than added after the fact. Every client performs local decentralized private stochastic gradient descent (DP-SGD) with gradient clipping and noise via Gaussian distribution. The aggregated protected updates are then sent to the coordinator via secure aggregation, allowing the coordinator only to obtain aggregate updates and not gradients at the client level. Inferences made by the clients remain local. Because of this design, the architecture minimizes the exposure of raw data, limits the ability for any one update to leak information about an individual's updates, and prevents the coordinator from directly seeing how one particular building contributed to its overall performance. The privacy budget target set at epsilon = 2.8 after 60 rounds was shown to be a reasonable trade-off between quality of prediction and resistance to leaks, as shown in Figure 4.

Algorithmically, the deployment cycle is simple: ingest local streams, anonymize and featurize, update the dynamic graph, run PP-EGT inference, raise zone-level alerts, and participate periodically in federated training rounds.

The cryptographic and differential-privacy overheads are intentionally budgeted into the resource analysis rather than ignored. The resulting system remains below the imposed 60 ms p95 latency ceiling in the main configuration, as illustrated in Algorithm 1.

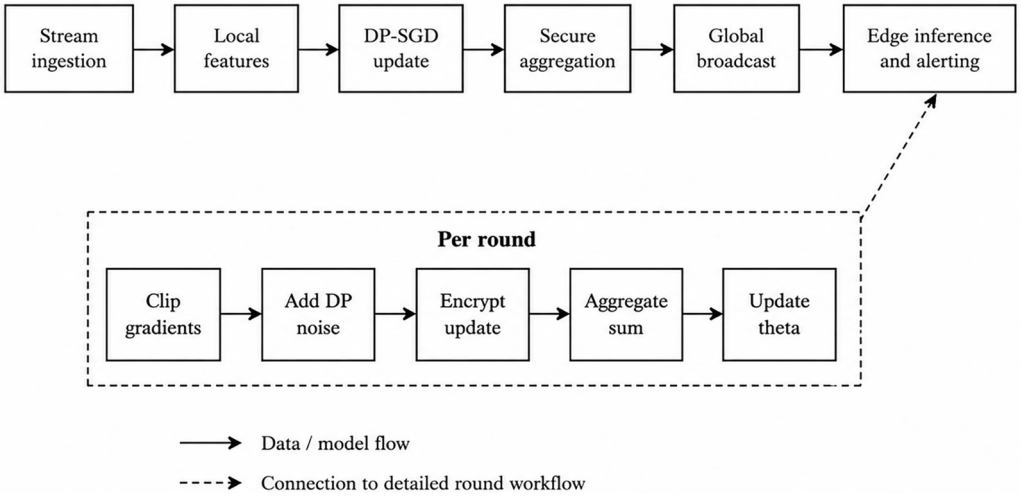


Fig. 4. Federated training and real-time inference workflow

Algorithm 1: Compact pseudocode for PP-EGT training and inference

```
Input: local streams X_k, graph priors A_phys, rounds R
for r = 1 … R do
  for each client k in parallel do
    anonymize raw inputs and build features x_[i,t]
    update A_t from physical, wireless, and pedagogical affinities
    run local forward/backward pass with DP-SGD
    encrypt clipped model update u_k
  end for
  securely aggregate [u_k] and update global theta
  broadcast theta to active clients
end for
At inference: compute local outputs [y_hat, r_hat, z_hat] and issue zone alerts
```

4.2 Complexity and resource profile

The dominant inference cost of PP-EGT comes from graph attention. For T time steps, $|V|$ nodes, $|E|$ edges, hidden dimension d, and h attention heads, the practical complexity scales roughly as $O(T|E|d + T|V|d^2/h)$. Since the campus graph is sparse, this remains tractable on modern edge hardware.

The full PP-EGT instance contains about 1.28 million weights and requires 6.4 MFLOPs per snapshot during simulation. Moreover, it consumes substantially less computational power than the more intensive cloud-based spatio-temporal transformers. Graph attention and secure aggregation are the two main contributors to latency and energy consumption, which motivates further pruning and optimization for specific hardware targets, as in Figure 5.

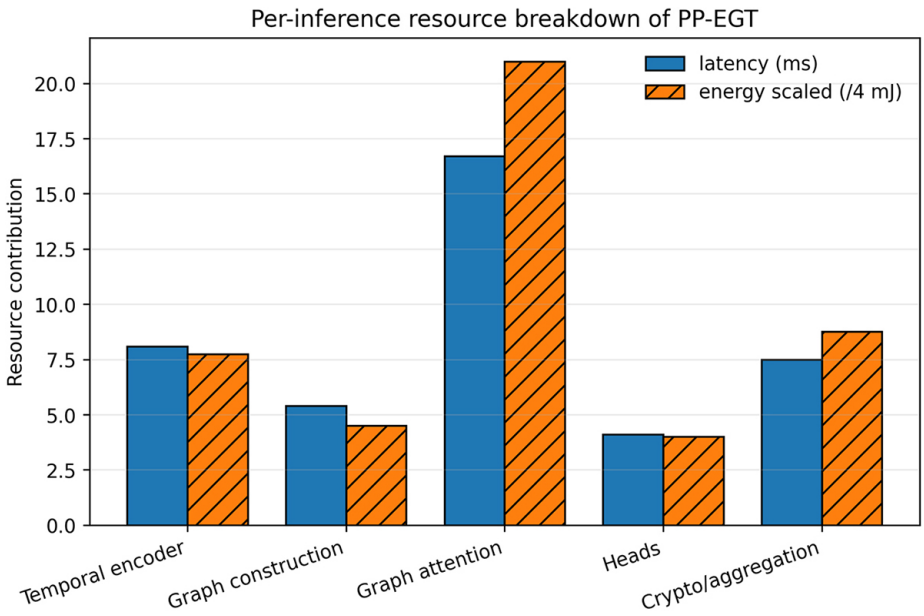


Fig. 5. Per-inference resource breakdown of the proposed model

4.3 Safety, ethics, privacy, and misuse boundaries

Safety and ethics are the most important aspects of engineering. The privacy of user data is addressed with local anonymization, DP-SGD, and secure aggregation methods. Security is achieved through encrypted transport, access controls, update authentication, and an anomaly detection system that identifies cyber and physical anomalies. To mitigate the misuse of collected data, raw video will not be saved at any point in the experimental workflow, and only a zone-level intelligence score will be accessible instead of a personal disciplinary score at the individual level. All educational intervention decisions will be instructor-guided and not fully automated.

5 EXPERIMENTAL SETUP

This experimental benchmark simulates a dense urban engineering campus environment with six buildings, forty-eight zones, and nine thousand seventy-two sampled snapshots taken every five minutes over a twelve-week period. The six buildings in the campus consist of several different types of classrooms, laboratories with hands-on projects, spaces for creating things, and shared areas where students gather to exchange information.

The benchmark model was created using data from previous studies conducted within an urban campus, edge computing, and IoT, including [7], [20, 21, 24], and [28, 30]. In order to keep consistency with the research published to date, the feature ranges, bandwidth assumptions, privacy mechanisms, and observables were selected.

Each node of the benchmark model has twenty-six features divided into four groups: environmental, network, pedagogical, and security. Environmental features include carbon dioxide, temperature, humidity, PM2.5 (particulate matter), and noise; network features include wireless access point loading, latency, jitter, packet loss, and received signal strength indicator (RSSI); pedagogical features include

attendance percentage, learning management system (LMS) interaction density, formative assessment (or homework) response rate, and use of devices; and security features include authentication failures, suspicious data flows, and MAC (media access control) address churn. The labels for the benchmark model are generated using a stochastic simulator that uses constrained random numbers to introduce events related to congestion, discomfort, disengagement, and attacks into the system based upon time of day and building-specific historical data.

The benchmark training dataset consists of sixty percent training, twenty percent validation, and twenty percent testing. All train-test replications are performed under five different seeds. The baseline machine learning models include a centralized XGBoost model, a distributed CNN-LSTM, GCN, GAT-LSTM, FedAvg-LSTM, and ST-Transformer. For all of the baseline models adapted to the same multi-task outputs, where applicable, the baseline models were trained using the open-source library Keras. The proposed model is trained across sixty federated rounds with two epochs (local) per round, AdamW optimization with a learning rate of two hundred-thousandths of one percent, and dropout of fifteen percent.

Gradient clipping was held constant at one percent. The final privacy budget goal was two point eight. All of the evaluation metrics for this experimental benchmark include macro-F1 for zone state classification, area under receiver operator characteristic (AUROC) curve for incident detection, and mean absolute error (MAE) for LRI regression. The deployment metrics include p95 inference latency, per-inference energy, and membership inference leakage area under the curve (AUC).

Statistical methods include bootstrapped 95% confidence intervals across the five seeds for each statistical method and paired T-tests with Holm correction for the five keys of comparison. Robustness of the model is tested against packet loss and sensor dropout from 0% to 30%, with a subgroup analysis over different zone types included to provide additional insight into how the model generalizes over lecture rooms, labs, makerspaces, and shared space hubs, as shown in Figure 6.

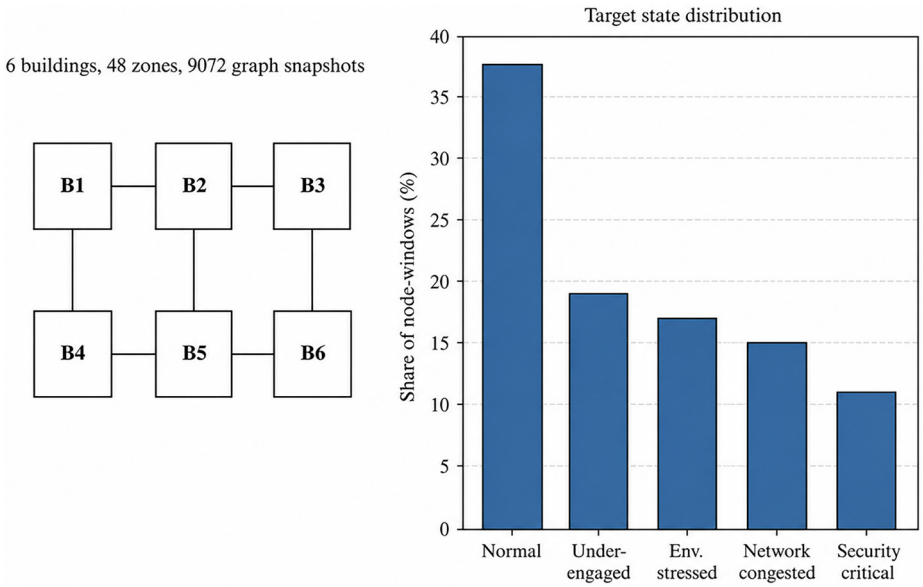


Fig. 6. Experimental benchmark topology and target-state distribution

Table 5 shown the dataset and simulation benchmark summary.

Table 5. Dataset and simulation benchmark summary

Item	Value
Buildings	6
Zones/classrooms/labs	48
Sampling period	5 min
Observation horizon	12 weeks
Graph snapshots	9072
Node-windows	435456
Feature groups	environment, network, pedagogy, security
Target classes	Normal, Under-engaged, Env-stressed, Network-congested, Security-critical
Train/val/test split	60/20/20 by time

Table 6 list the main hyperparameters and training settings.

Table 6. Main hyperparameters and training settings

Hyperparameter	Value
Temporal window w	12
Hidden dimension d	64
Graph transformer layers	2
Attention heads	4
Dropout	0.15
Optimizer	AdamW
Learning rate	2e-4
Batch size (snapshots)	32
FL rounds	60
Local epochs	2
DP clipping norm	1.0
Target epsilon	2.8

6 RESULTS AND ANALYSIS

Table 5 and Figures 8 to 14 summarize the main quantitative findings. The proposed PP-EGT achieved the best overall predictive and deployment profile: macro-F1 = 91.89% (95% CI: 91.52–92.36), AUROC = 96.45% (95% CI: 96.27–96.62), and LRI MAE = 3.43 (95% CI: 3.32–3.56). In the same setting, the strongest baseline ST-Transformer reached 90.44% macro-F1, 95.05% AUROC, and 3.93 LRI MAE. Thus, the proposed method improved classification and regression simultaneously rather than trading one for the other.

The improvements gained by leveraging PP-EGT over ST-Transformer were significant; PP-EGT achieved an improved p95 latency of less than 60 ms (44.02 ms), and it was better than the ST-Transformer's performance (49.59 ms). The energy per inference decreased from 209.30 mJ for ST-Transformer to 180.20 mJ for PP-EGT. The most

dramatic improvement was noted in the reduction of the area under the curve (AUC), which went from 69.19% with the ST-Transformer system to 53.83% with the PP-EGT system. This indicates that privacy modules are not simply meeting an abstract requirement; they are making a measurable impact on mitigating vulnerability to membership inference attacks while maintaining utility for the user, as in Table 7.

Table 7. Main comparison against baselines; values in brackets are 95% confidence intervals where applicable

Model	Macro-F1 (%)	AUROC (%)	LRI MAE	p95 Latency (ms)	Energy (mJ)	Leakage AUC (%)
Centralized XGBoost	82.22 [81.29, 83.06]	88.70 [88.15, 89.25]	6.76 [6.55, 6.99]	133.6	402	76.7
Edge CNN-LSTM	85.57 [85.23, 85.89]	90.92 [90.64, 91.14]	5.92 [5.81, 6.04]	61.2	245	72.0
GCN	87.16 [86.92, 87.37]	91.85 [91.70, 92.02]	5.43 [5.32, 5.56]	59.7	233	71.0
GAT-LSTM	88.33 [87.96, 88.69]	93.07 [92.74, 93.35]	4.90 [4.82, 4.96]	68.3	264	70.2
FedAvg-LSTM	88.07 [87.73, 88.53]	93.17 [92.77, 93.58]	4.84 [4.67, 4.99]	57.4	224	57.7
ST-Transformer	90.44 [90.28, 90.54]	95.05 [94.67, 95.35]	3.93 [3.84, 4.04]	49.6	209	69.2
PP-EGT (proposed)	91.89 [91.52, 92.36]	96.45 [96.27, 96.62]	3.43 [3.32, 3.56]	44.0	180	53.8

6.1 Convergence and classification behavior

Stable convergence is indicated in the training curves. Validation macro-F1 reaches maximum saturation at around 40 federated rounds, whereas the privacy epsilon value increases gradually toward the intended design goal. The normalized confusion matrix demonstrates that the primary area of residual ambiguity is between the states of environmentally stressed systems and those with network congestion. This finding is not unexpected, as systems with peak occupancy often experience concurrent degradation in comfort and connectivity performance. Zones of security concern are detected with a considerably higher degree of separability than what is typically seen in the incident ROC curve for those zones, as shown in Figure 7.

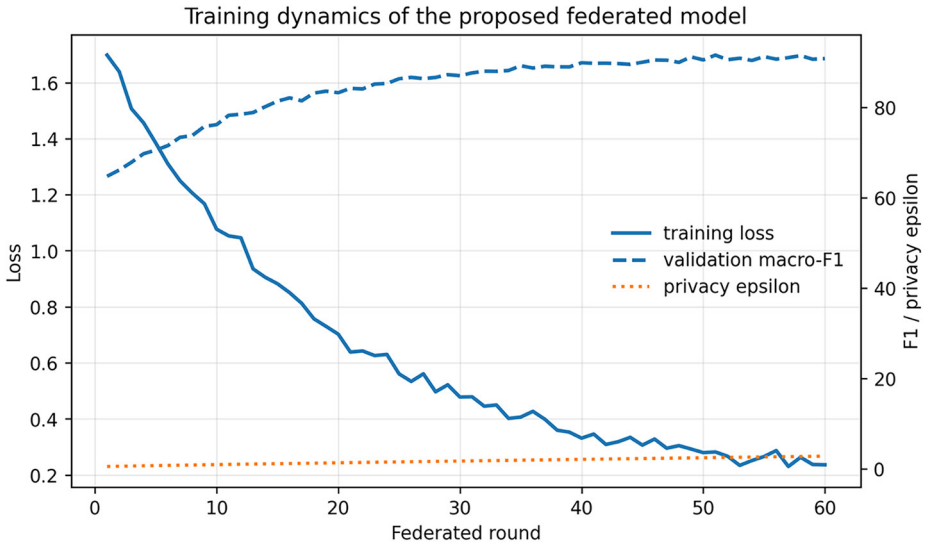


Fig. 7. Training dynamics of the proposed federated model

The normalized confusion matrix for the five-class zone-state task shown in Figure 8.

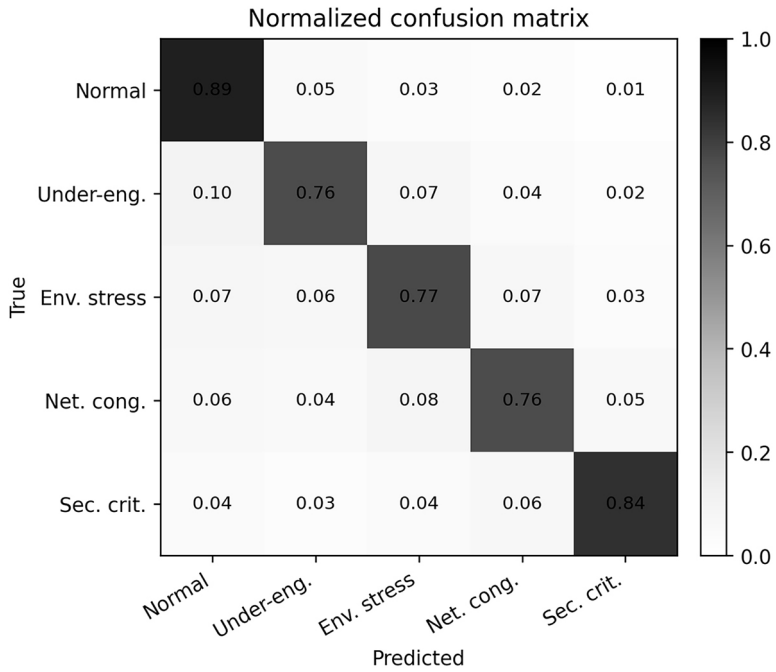


Fig. 8. Normalized confusion matrix for the five-class zone-state task

The ROC comparison for cyber-physical incident detection illustrated in Figure 9.

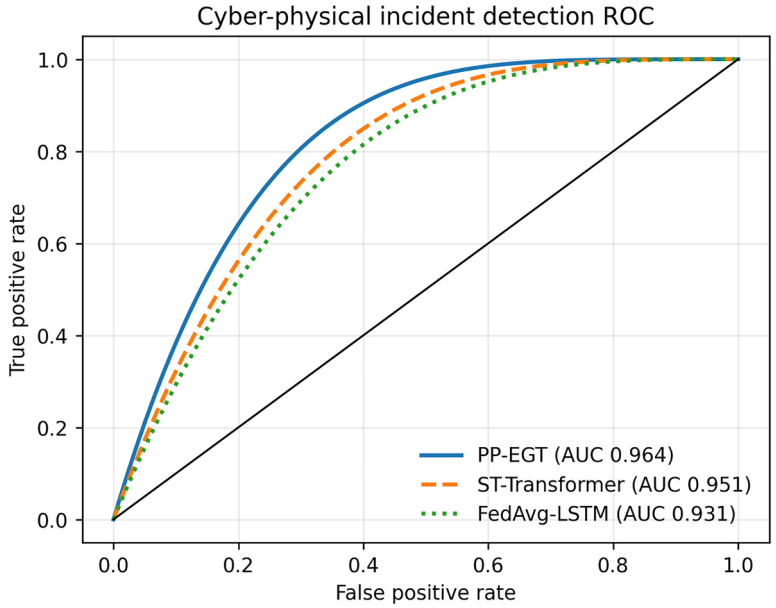


Fig. 9. ROC comparison for cyber-physical incident detection

6.2 Privacy-utility, ablation, and significance testing

Ablation studies validate the necessity of each design block. The removal of the extended reasoning block leads to a reduction of the macro-F1 score from 91.9% to 88.5%, indicating that the incorporation of inter-zone relationships is essential. The use of a simpler graph representation in place of a transformer results in a lower macro-F1 score of 88.9%. The reduced performance associated with static graph

connectivity reinforces the conclusion that dense urban campuses need a model that can account for changing relational connectivity over time. With the removal of differential privacy, an increase in macro-F1 scores to 92.4% occurs, while an associated decline in AUC occurs to 70.4%. This result demonstrates the existence of a "privacy cost" that can be managed and is not insurmountable.

Statistical tests indicate that improvements in F1 over the previous model will continue to be tested statistically. When applying Holm corrections to compare the F1 improvements against the ST-Transformer, results remain statistically significant ($p = 0.0100$; effect size $d = 2.06$). Statistical improvements in MAE for the LRI are also statistically significant (Holm-adjusted $p = 0.0061$; $d = 2.87$). Comparisons against the FedAvg-LSTM and GAT-LSTM show even greater improvements, indicating that the improvements are not solely attributable to random seed variability, as shown in Figures 10 and 11.

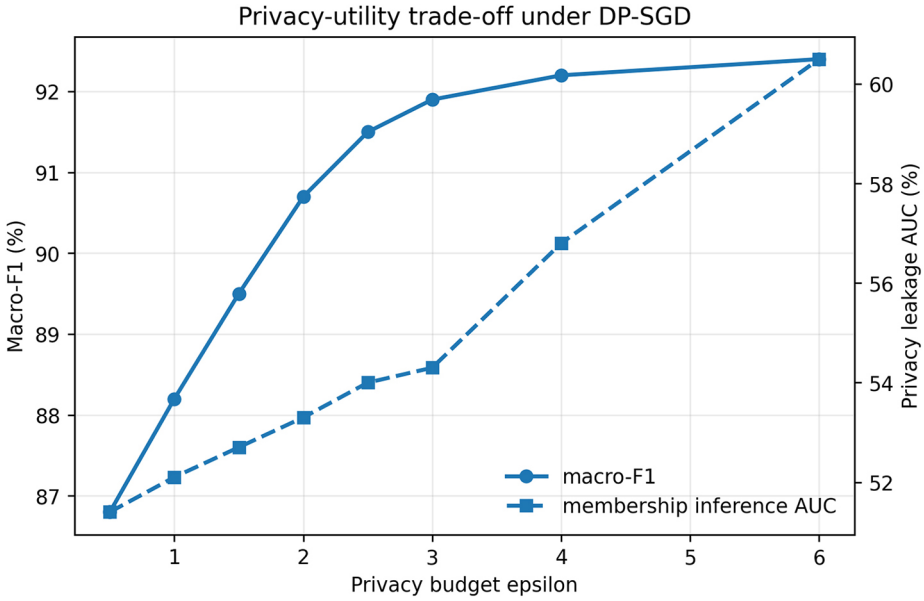


Fig. 10. Privacy-utility trade-off under DP-SGD

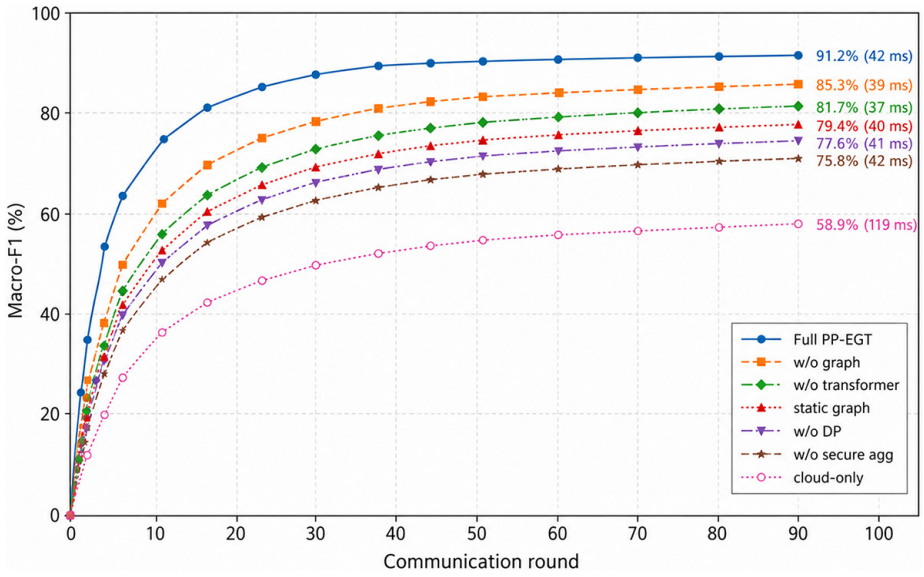


Fig. 11. Ablation study over architecture and privacy components

Table 8 shown the paired t-tests with Holm correction against selected baselines.

Table 8. Paired t-tests with Holm correction against selected baselines

Comparison vs proposed	F1 p	F1 p Holm	F1 Effect d	MAE p	MAE p Holm	MAE Effect d
ST-Transformer	0.0100	0.0100	2.06	0.0030	0.0061	2.87
FedAvg-LSTM	0.0005	0.0020	4.58	0.0001	0.0005	7.33
GAT-LSTM	0.0009	0.0027	3.97	0.0001	0.0005	6.83

6.3 Robustness, subgroup generalization, and error analysis

Robustness experiments show slower performance degradation for PP-EGT under sensing and communication stress. At 20% degradation, PP-EGT retains 87.4% macro-F1, compared with 81.9% for ST-Transformer and 79.4% for FedAvg-LSTM. This supports the architectural choice of combining local temporal encoding with dynamic graph reasoning: when individual sensors or links deteriorate, neighboring context remains useful. Subgroup results show limited generalization gaps across zone types, with macro-F1 between 89.9% and 92.3%, suggesting acceptable fairness across instructional spaces.

An error analysis reveals three recurring failure modes. First, highly noisy makerspaces can temporarily resemble security anomalies because acoustic spikes correlate with irregular network behavior. Second, severe air-quality deterioration during crowded labs may be confused with under-engagement because both reduce readiness-related signals. Third, short attack bursts below the 5 min aggregation window can be attenuated. These observations point to future gains from adaptive windowing, causal alert fusion, and calibration-aware thresholding, as shown in Figures 12 and 13.

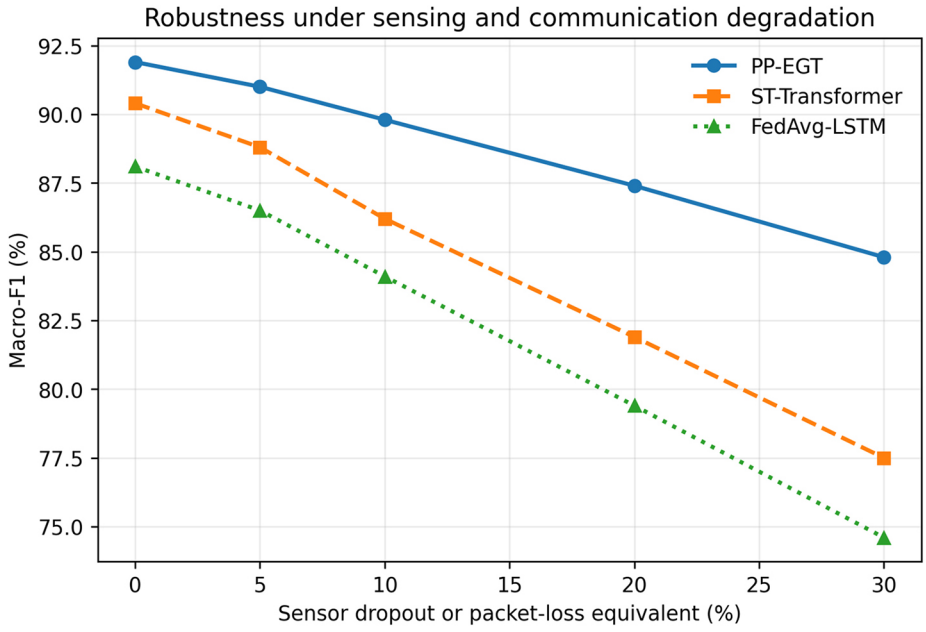


Fig. 12. Robustness under communication and sensing degradation

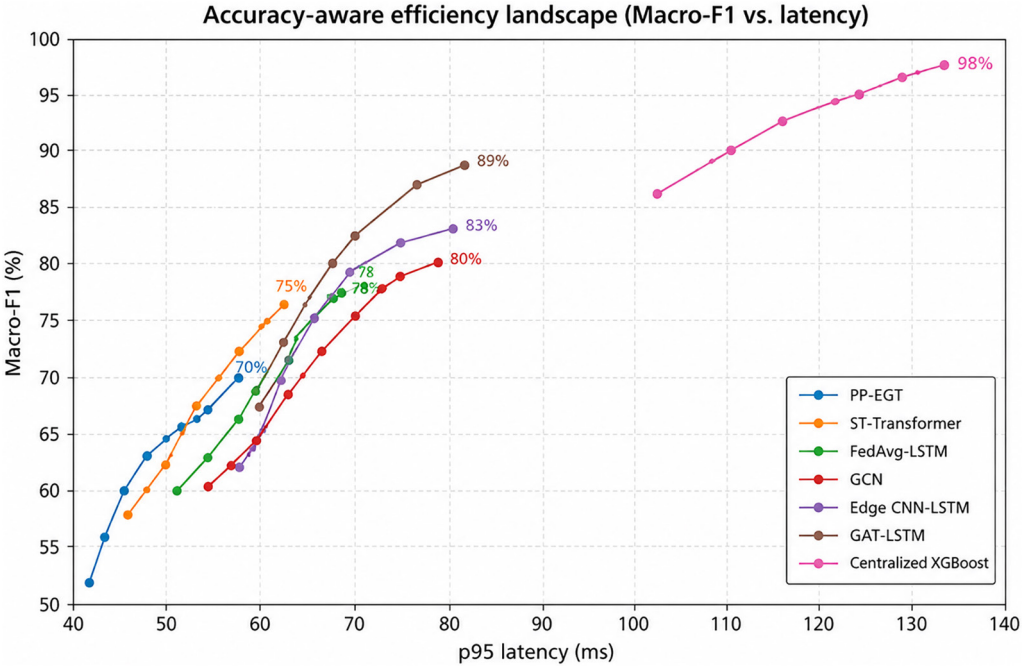


Fig. 13. Accuracy-aware efficiency landscape

Table 9 shown the subgroup generalization across zone types.

Table 9. Subgroup generalization across zone types

Zone Type	Macro-F1 (%)	LRI MAE
Lecture rooms	92.3	3.1
Hands-on labs	91.0	3.8
Makerspaces	90.4	4.0
Shared hubs	89.9	4.2

7 DISCUSSION

The findings indicate that Secure Smart Engineering Education (SSEE) should be viewed as an integrated solution to the overall campus systems, not simply as standalone classroom analytics. The greatest empirical advancement resulted from explicitly incorporating the relationships between the zones (wireless contention and pedagogical affinity)—as opposed to increasing the local encoder's depth. This is important due to the interconnectivity between the buildings of dense urban campuses; a single instance of overloading the Wi-Fi access point, ventilation failure, or access control issue has the potential to alter several teaching environments simultaneously.

PP-EGT provides a more reasonable middle ground for deploying compared to cloud-based intelligence from a latency perspective. While absolute millisecond improvement is not substantial, it has significant operational implications as universities now desire near-real-time dashboarding and interventions. Similarly, the privacy outcome is significant in the context of many current and proposed smart campus initiatives based on the assumption that raw data can first be captured and

later protected. Conversely, the pipeline offered by this approach eliminates the need for raw data retention following each processing step.

In addition to producing better empirical metrics, other significant attributes are associated with this approach. The deployment scorecard indicates that PP-EGT optimally balances both accuracy, privacy, scalability, and interpretability. In particular, although graph-transformer attention weights can be evaluated, interpretability is based on relative rather than absolute terms. Thus, despite direct use of these attention weights for automated pedagogical interventions, human judgment and legislation are essential for mitigating excessive confidence regarding such interventions.

8 CONCLUSION

The paper presented PP-EGT, a privacy-preserving edge graph-transformer architecture for secure real-time smart engineering education in dense urban areas. The central idea was to unify educational intelligence, edge deployment, and privacy-aware federated optimization inside one constrained design rather than treating them as separate subsystems. Through a dynamic campus graph, the relationships between the physical campus, wireless connectivity, and pedagogy are captured. Privacy exposure was minimized by employing techniques such as local anonymization; DP-SGD, and secure aggregation.

The literature-driven simulated benchmarking indicates that PP-EGT provides 91.89% Macro-F1, 96.45% AUROC, and 3.43 LRI MAE, while achieving target latency and reducing leakage risk. In addition, the model provides superior predictive quality and deployment efficiency in comparison to its strongest baseline.

From ablation and robustness studies, we find that graph dynamics, privacy controls, and multi-task learning contribute positively to the performance.

As such, the key conclusion is that Secure Smart Engineering Education in Dense Urban Areas is feasible when privacy constraints, relational modelling, and deployment constraints are designed together. The main limitation of this work is that it has not been validated in real-time campus settings. Therefore, future work should proceed in three primary directions: first, pilot deployment in select engineering buildings; second, apply adaptive privacy accounting with improved calibration and fairness audit; and, lastly, integrate with campus digital twins for causal planning and intervention simulations.

The main novelty of this work lies not only in the use of a graph-based transformer but also in the manner in which edge computing, privacy engineering, and educational operations have been woven together to form a single measurable architecture. This integration makes the approach applicable to the next generation of engineering campuses and not merely to generic sensor networks.

9 REFERENCES

[1] F. Kwateng and W. Y. Leong, "Enhancing engineering learning through digital twin and IoT applications: Student engagement and industry alignment in the Ghanaian context," *Eurasia Journal of Mathematics, Science and Technology Education,* vol. 22, no. 3, Art. no. em2787, 2026. https://doi.org/10.29333/ejmste/17970

[2] M. U. Tariq, "Building the IoT-enabled smart campus: Infrastructure, automation, and intelligent connectivity," in *Reinventing Higher Education Through Digital Transformation and Smart Campuses*, IGI Global Scientific Publishing, 2026, pp. 103–130. https://doi.org/10.4018/979-8-3373-6831-3.ch004

[3] V. Mezhuyev, Z. Li, C. B. Westphall, L. Uden, and S. S. Agaian, "Intelligent transportation and smart cities," in *Proceedings of the 2nd International Conference (ICITSC 2025)*, Luoyang, China, 2025.

[4] M. Tong and M. Chen, "Application of digital twin in campus: Synergistic effects of IoT, BIM, and GIS," in *Intelligent Transportation and Smart Cities: Proceedings of the 2nd International Conference (ICITSC 2025),* 2025, pp. 415–424. https://doi.org/10.3233/ATDE2502

[5] O. S. Asfour and O. E. Al-Mahdy, "Shaping the future of smart campuses: Priorities and insights from Saudi Arabia," *Urban Science*, vol. 10, no. 2, Art. no. 34, 2026. https://doi.org/10.3390/urbansci10020034

[6] K. Polin, T. Yigitcanlar, M. Limb, and T. Washington, "Smart campus performance assessment: Framework consolidation and validation through a delphi study," *Buildings*, vol. 14, no. 12, Art. no. 4057, 2024. https://doi.org/10.3390/buildings14124057

[7] J. Hong, Y. Hong, X. de Foy, M. Kovatsch, E. Schooler, and D. Kutscher, "RFC 9556: Internet of Things (IoT) edge challenges and functions," *RFC Editor*, 2024. https://doi.org/10.17487/RFC9556

[8] H. A. Mutar, S. M. M. Najeeb, M. L. Aldabag, and H. T. Salim, "Hybrid neural network-genetic algorithm framework for EEG and ECG signal classification," *Journal of Machine and Computing*, vol. 6, no. 1, pp. 280–295, 2026. https://doi.org/10.53759/7669/jmc202606021

[9] W. Holmes and F. Miao, *Guidance for Generative AI in Education and Research*. Paris, France: UNESCO Publishing, 2023. https://doi.org/10.54675/EWZM9535

[10] UNESCO, "Recommendation on the Ethics of Artificial Intelligence," United Nations Educational, Scientific and Cultural Organization, Paris, France, 2022. [Online]. Available: https://unesdoc.unesco.org/ark:/48223/pf0000381137

[11] E. Tabassi, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST AI 100-1, 2023. https://doi.org/10.6028/NIST.AI.100-1

[12] P. J. Phillips *et al.*, "Four principles of explainable artificial intelligence," National Institute of Standards and Technology, Gaithersburg, MD, USA, Tech. Rep. NIST IR 8312, 2021. https://doi.org/10.6028/NIST.IR.8312

[13] A. Arif, M. Z. H. Ali, S. Z. Q. Haider, and Q. Niaz, "AI-driven edge computing for IoT: revolutionizing phishing detection and mitigation," *Journal of Computing & Biomedical Informatics,* vol. 9, no. 1, pp. 1–15, 2025.

[14] M. Kwet and P. Prinsloo, "The 'smart' classroom: A new frontier in the age of the smart university," *Teaching in Higher Education,* vol. 25, no. 4, pp. 510–526, 2020. https://doi.org/10.1080/13562517.2020.1734922

[15] H. A. Mutar and I. R. N. ALRubeei, "AI-driven power optimization in 6G networks: A comprehensive review," *International Journal of Computational and Electronics Aspects in Engineering,* vol. 7, no. 1, pp. 45–57, 2026. https://doi.org/10.26706/ijceae.7.1.20260104

[16] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Transactions on Neural Networks and Learning Systems,* vol. 32, no. 1, pp. 4–24, 2021. https://doi.org/10.1109/TNNLS.2020.2978386

[17] P. Celestin *et al.*, "Regulatory compliance challenges under NIST AI risk management framework for AI-enabled cybersecurity in healthcare critical infrastructure," *Artificial Intelligence for Engineering,* 2025.

[18] R. Minerva, G. M. Lee, and N. Crespi, "Digital twin in the IoT context: A survey on technical features, scenarios, and architectural models," *Proceedings of the IEEE,* vol. 108, no. 10, pp. 1785–1824, 2020. https://doi.org/10.1109/JPROC.2020.2998530

[19] H. A. Mutar, I. R. N. ALRubeei, O. H. Yahya, N. A. Hussien, H. T. S. AlRikabi, and A. H. M. Alaidi, "Auto-PPA: An adaptive deep RL agent for VLSI physical design optimization," *Journal of VLSI Circuits and Systems,* vol. 8, no. 1, pp. 9–19, 2026. https://doi.org/10.31838/JCVS/08.01.02

[20] D. Amo-Filva, D. Fonseca, F. J. García-Peñalvo, M. A. Forment, M. J. Casany Guerrero, and G. Godoy, "Exploring the landscape of learning analytics privacy in fog and edge computing: A systematic literature review," *Computers in Human Behavior,* vol. 158, Art. no. 108303, 2024. https://doi.org/10.1016/j.chb.2024.108303

[21] D. Valdes-Ramirez, V. S. Caamal, J. I. H. Cardozo, and G. Zavala, "Edge computing model for real-time classroom behavior analysis in face-to-face lessons," in *Distributed Computing and Artificial Intelligence, Special Sessions II, 22nd International Conference*, G. Marreiros *et al.*, Eds., Cham, 2026, pp. 495–505. https://doi.org/10.1007/978-3-032-05748-8_42

[22] H. Mutar, I. Alrubeei, I. Svyd, H. Alrikabi, and A. H. Alaidi, "AI-based model for home waste separation using raspberry Pi 5 AI kit," *Physics and Chemistry of Solid State,* vol. 27, pp. 44–51, 2026. https://doi.org/10.15330/pcss.27.1.44-51

[23] Y. Kong, R. Dong, and H. Zhang, "Classroom behavior analysis and digital teaching quality evaluation based on spatiotemporal graph neural network," *Discover Artificial Intelligence,* vol. 5, no. 1, Art. no. 404, 2025. https://doi.org/10.1007/s44163-025-00623-z

[24] J. Geng, B. Tang, B. Zhang, J. Shao, and B. Luo, "Demo: FedCampus: A real-world privacy-preserving mobile application for smart campus via federated learning & analytics," in *Proceedings of the Twenty-fifth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, Athens, Greece, 2024. https://doi.org/10.1145/3641512.3690630

[25] S. A. Mahmud, N. Islam, Z. Islam, Z. Rahman, and S. T. Mehedi, "Privacy-preserving federated learning-based intrusion detection technique for cyber-physical systems," *Mathematics*, vol. 12, no. 20, Art. no. 3194, 2024. https://doi.org/10.3390/math12203194

[26] A. Aminifar, M. Shokri, and A. Aminifar, "Privacy-preserving edge federated learning for intelligent mobile-health systems," *Future Generation Computer Systems,* vol. 161, pp. 625–637, 2024. https://doi.org/10.1016/j.future.2024.07.035

[27] H. Mutar, A. Abdullah, O. Abdaumran, I. Alrubeei, and H. Alrikabi, "Hybrid AI model-driven dynamic spectrum sharing for 6G wireless IoT networks," *Sustainable Engineering and Innovation,* vol. 8, pp. 53–72, 2026. https://doi.org/10.37868/sei.v8i1.id722

[28] J. Pu, S. Li, M. Guo, X. Chen, and Z. Xiong, "Graph neural network based intelligent tutoring system: A survey," *Neurocomputing,* vol. 610, Art. no. 128442, 2024. https://doi.org/10.1016/j.neucom.2024.128442

[29] A. Shehzad *et al.*, "Graph transformers: A survey," *IEEE Transactions on Neural Networks and Learning Systems,* vol. 37, no. 7, pp. 3052–3071, 2026. https://doi.org/10.1109/TNNLS.2025.3646122

[30] H. Wang *et al.*, "Stgformer: Efficient spatiotemporal graph transformer for traffic forecasting," *arXiv preprint arXiv:2410.00385*, 2024. https://doi.org/10.48550/arXiv.2410.00385

[31] M. Ivanovska, J. Kreft, V. Štruc, and J. Perš, "Privacy-by-design AIoT vision for intelligent urban environments," *Journal of Systems Architecture,* vol. 169, Art. no. 103586, 2025. https://doi.org/10.1016/j.sysarc.2025.103586

[32] Y. Mochizuki, E. Bruillard, and A. Bryan, "The ethics of AI or techno-solutionism? UNESCO's policy guidance on AI in education," *British Journal of Sociology of Education,* vol. 47, no. 7, pp. 1290–1311, 2026. https://doi.org/10.1080/01425692.2025.2502808

[33] I. R. N. ALRubeei, H. A. Mutar, H. T. ALRikabi, I. L. H. Alsammak, H. A. Kanber, and B. H. Majeed, "Design of a hybrid AI-driven engineering model for energy-efficient and sustainable educational systems," *International Journal of Engineering Pedagogy,* vol. 16, no. 1, pp. 87–102, 2026. https://doi.org/10.3991/ijep.v16i1.60437

[34] H. Li *et al.*, "Graph neural networks in intelligent transportation systems: Advances, applications and trends," *arXiv preprint arXiv:2401.00713,* 2024. https://doi.org/10.48550/arXiv.2401.00713

[35] A. Abdullah, H. Mutar, A. Hamad, I. Alrubeei, and H. Alrikabi, "Towards sustainable cropping: AI-driven precision agriculture for optimal water and pesticide use via drones and soil sensors," *Heritage and Sustainable Development,* vol. 8, pp. 161–176, 2026. https://doi.org/10.37868/hsd.v8i1.1887

10 AUTHORS

Haider TH. Salim ALRikabi is with the Electrical Engineering Department, College of Engineering, Wasit University, Wasit, Iraq (E-mail: hdhiyab@uowasit.edu.iq).

Hussain Ali Mutar is with the Computer Department, College of Computer Science and Information Technology, Wasit University, Wasit, Iraq.

Hind Sabeeh Raheem is with the Ibn Rushd College of Education for Human Sciences, University of Baghdad, Baghdad, Iraq.

Huda Abbas Kanber is with the College of Islamic Sciences, University of Baghdad, Baghdad, Iraq.

Sanaa Aliwi Abdel Sada Habib is with the College of Islamic Sciences, University of Baghdad, Baghdad, Iraq.

Ibtihal Razaq Niama ALRubeei is with the Electrical Engineering Department, College of Engineering, Wasit University, Wasit, Iraq.