SSL-FSL Framework For Detecting Unseen IoT Attacks Under Limited Labeled Data

Amer A.Abualhassan*, Louai Al-Awami*§, Mosab Hamdan‡
*Computer Engineering Department, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
§Interdisciplinary Research Center for Intelligent Secure Systems, KFUPM, Saudi Arabia
‡School of Computing, National College of Ireland, Dublin, Ireland
Email: G202214400@kfupm.edu.sa, louai@kfupm.edu.sa, mosab.mohamed@ncirl.ie

***Abstract*—Traditional intrusion detection system solutions and modern machine learning (ML) and deep learning (DL)-based approaches depend on large labeled datasets, which makes them less effective in detecting rare attacks in Internet of Things (IoT) networks. This paper addresses these limitations by proposing a self-supervised few-shot learning (SSL-FSL) intrusion detection framework that enables data-efficient learning and adaptive recognition of unseen attack types using only a small number of labeled samples. The model will be trained and tested on the Edge-IIoTset dataset as well as on unseen attack categories from the CICIoT2023 dataset. The traffic data will be represented as images to allow the model to learn spatial feature patterns rather than relying on class-specific characteristics. This representation, combined with self-supervised pre-training and prototype-based few-shot adaptation, enables the extraction of robust and transferable embeddings that capture intrinsic relationships within the data. Experimental results demonstrate that the proposed SSL-FSL framework achieves strong unseen attack detection performance and achieves an accuracy of 69.76% on the CICIoT2023 dataset without any fine-tuning, highlighting its ability to generalize to novel attack types. Overall, the proposed framework provides a practical, data-efficient, and generalizable solution for adaptive intrusion detection in heterogeneous IoT environments.**

***Index Terms*—Internet of Things, intrusion detection system, few-shot learning, self-supervised learning, limited labeled data, deep learning, network security**

I. Introduction

Internet of Things (IoT) technology has rapidly expanded to many different areas. However, as the number of IoT devices expands, the attack surface for those devices has expanded as well. Ensuring that these devices are protected from such threats is a critical challenge in the expansion of the IoT technology.

Intrusion Detection Dystem (IDS) have traditionally been used to identify malicious activities in networks. Recent advances in ML and DL have improved detection accuracy and enabled the identification of complex attack patterns [1]–[3]. Despite these advancements, most ML- and DL-based approaches rely heavily on large labeled datasets, which are difficult to obtain in IoT environments due to privacy constraints and the rarity of many attack types [4]. In addition, IoT traffic is often event-driven, leading to limited data availability during low activity periods, further reducing traffic generation. Consequently, ML- and DL-models struggle to generalize to unseen or emerging attacks.

978-3-903176-82-9

The scarcity and imbalance of labeled data are significant challenges in IoT IDS. The nature of IoT environments to be dynamic and ever-changing makes the supervised models trained on a fixed dataset ineffective in real IoT environments. These limitations have been highlighted in prior work [3], where the dependency on data availability and resource constraints significantly impacted model performance. In parallel, recent work has explored lightweight IDS for resource-constrained IIoT, using model compression for efficient and effective deployment [5].

In order to overcome these limitations, we propose a framework that integrates these approaches in order to enable the model to recognize new types of attacks in the IoT systems while requiring minimal labeled data for training. Therefore, the goal of this research is to propose a framework that learns representations from unlabeled data and learns to recognize new attack types in the IoT with minimal supervision. Furthermore, the framework will be evaluated on both seen and unseen IoT attack datasets. The main contributions of this paper are as follows:

- We propose an SSL-FSL framework for IoT intrusion detection to tackle the challenges of limited labeled dataset, class imbalance, and dynamic behavior of attacks on IoT networks.
- By leveraging an existing representation of the network traffic as images, we can apply convolutional neural networks to extract features that may be indicative of anomalous behavior and learn generalizable network embeddings
- We extend the application of this transformation by converting the CICIoT2023 dataset using the same image-based strategy, enabling a consistent and fair evaluation on completely unseen attack categories.
- We design a learning framework that jointly optimizes the self-supervised pre-training and few-shot adaptation stages to recognize previously unseen attack classes with only a few labeled samples.
- We evaluate the proposed framework on both seen (Edge-IIoTset) and unseen (CICIoT2023) datasets without any fine-tuning, demonstrating strong generalization capability in cross-dataset intrusion detection scenarios.

The remainder of this paper is organised as follows:

Section II reviews the related work. Section III presents the Methodology. Section IV describes the experimental setup. Section V presents the results of the experiments. Section VI discussion and Finally, the conclusions of this work.

II. Related Work

Recent works have begun to address data scarcity and class imbalance in network intrusion detection by combining FSL with SSL. For example, Dina et al.*et al.* [6] proposes the FS3 framework, which operates in three phases to secure IoT networks. In the first phase, a self-supervised objective learns latent patterns from unlabeled network traffic, producing robust feature representations. In the second phase, few-shot learning with contrastive training (e.g., triplet loss) is performed to adapt the model to new attack classes using only a handful of labeled examples. This method also helps balance the highly imbalanced class distribution by improving the modelâs ability to discriminate minority classes. In the last phase, a k-NN method is used where the samples from the majority classes are down-sampled to rebalance the training data. The results of this approach show that using only 20% of the labeled data, FS3 outperformed the fully supervised method by 42% in precision and 43% in F1-score, thus significantly improving the detection of minority attack classes.

Similarly, Atitallah *et al.* [7] propose a hybrid SSL-FSL approach to strengthen the detection of network intrusion attacks using a Deep InfoMax and a prototypical network to achieve 99% detection accuracy. In a complementary direction, Xu *et al.* [8] introduces an enhanced few-shot network IDS based on a transformer architecture with self-attention and iterative meta-learning refinement. The proposed model employs embedding, encoding, and classification layers, where self-attention mechanisms capture long-range dependencies in network traffic and positional encoding preserves temporal information. The model is trained using a meta-learning approach where the model learns on a set of base intrusion detection tasks and is adapted to a few samples of new intrusion detection classes. The adaptation is performed iteratively to allow for more effective utilization of the few available training samples. The proposed model achieves detection rates close to 99% for the few-shot detection of network intrusion using as few as 10 training samples. Furthermore, each iteration of the model introduces only a minimal computational overhead of approximately 1.9 ms while improving the detection accuracy of the intrusion detection system.

Beyond the cybersecurity domain, the integration of SSL and FSL has been widely studied in computer vision and other machine learning fields. For instance, Gidaris *et al.* [9] and Lim *et al.* [10] demonstrate that SSL significantly enhances feature representation quality in few-shot scenarios. Their methods were evaluated on standard benchmarks such as MiniImageNet, tiered-MiniImageNet, and CIFAR-FS, which are widely used to assess generalization performance under limited labeled data. The success of SSL-enhanced FSL methods on these datasets provides strong evidence for their applicability in more complex and domain-specific problems such as IoT intrusion detection.

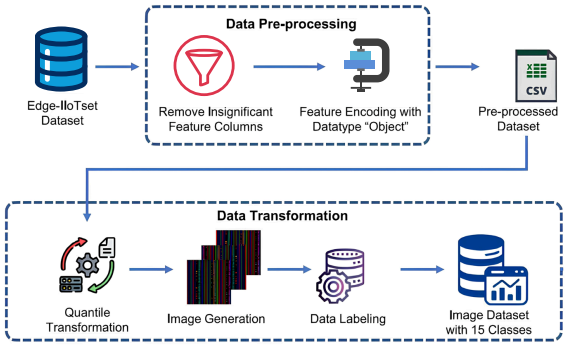

Fig. 1: Data preprocessing and transformation pipeline. Raw IoT traffic data is cleaned, normalized, and converted into image representations for SSL–FSL-based intrusion detection.

Furthermore, Latif *et al.* [11] introduced a method that effectively transforms IoT network traffic data into image form, which could then be fed into CNN models to detect network intrusions. The Edge-IIoTset dataset is transformed from CSV files to images, allow spatial feature extraction and improving detection performance.

In this study, we adopt the transformed dataset provided by Latif *et al.* [11], leveraging its image-based representation to benefit from advanced deep learning models and enhance feature learning under limited data conditions.

Inspired by these advances, this work proposes a metric-based few-shot intrusion detection framework, termed *SSL-FSL*, designed to effectively classify cyber-attacks under extreme data scarcity conditions.

III. Methodology

This section presents the design and implementation of the proposed framework. The framework is designed to preprocess data, transform tabular data to image form, and apply SSL and FSL techniques to enable the detection of rare and previously unseen attack classes with limited data.

A. Data Preprocessing and Feature Transformation

As illustrated in Fig. 1, raw IoT traffic data is first subjected to a preprocessing stage, followed by a transformation pipeline prior to model training. Initially, non-informative features are removed, and categorical attributes are encoded to generate a structured dataset suitable for learning.

For stable training, the feature values are normalized using quantile transformation before being transformed into multi-channel images as described by Latif *et al.* [11]. The same transformation is applied to the CICIoT2023 dataset for evaluating the model. As the CICIoT2023 dataset has 46 features as compared to the 95 features of the Edge-IIoTset dataset, the size of the images is set to $46 \times 46 \times 3$ RGB images are generated to maintain consistent representation.

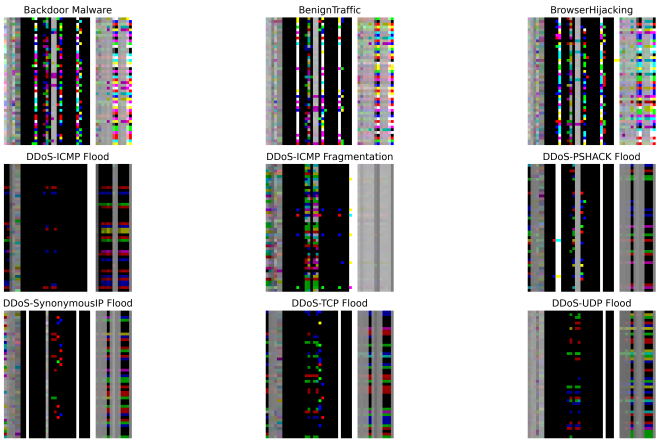

Fig. 2: Sample's image representations generated from IoT traffic data after tabular-to-image transformation. Each image corresponds to a traffic class (Benign Traffic or Attacks).

Each generated image is assigned a label corresponding to the dominant class within the grouped samples, ensuring consistency between representation and class identity.

As shown in Fig. 2, the tabular to image transformation produces visually distinct patterns for different traffic categories, including Benign Traffic and various attack types such as Backdoor Malware, Browser Hijacking, DDoS-ICMP Flood, DDoS-ICMP Fragmentation, DDoS-PSHACK Flood, DDoS-SynonymousIP Flood, DDoS-TCP Flood, and DDoS-UDP Flood obtained from the CICIoT2023 dataset after tabular-to-image transformation.

B. SSL–FSL Intrusion Detection Framework

To address the challenges of detecting rare and previously unseen threats, the proposed framework employs a hybrid SSL–FSL model that enables adaptive learning under limited labeled data conditions.

As illustrated in Fig. 3, the framework follows a two-stage learning paradigm that combines self-supervised representation learning with prototype-based few-shot classification. In the first stage, the model learns general feature representations from unlabeled data, while in the second stage, these representations are adapted for classification using a limited number of labeled samples.

The SSL–FSL model utilizes a shared encoder architecture, as shown in Fig. 4. This encoder is first trained in a self-supervised manner to capture intrinsic traffic patterns and is subsequently adapted using few-shot learning for classification. The learning process consists of two main stages:

- **Self-Supervised Pre-Training:** The encoder learns rotation-invariant and structurally meaningful representations from unlabeled data.
- **Few-Shot Fine-Tuning:** Using a small labeled support set, the model constructs class prototypes and classifies query samples based on similarity. A secondary self-supervised objective is incorporated to improve stability and reduce overfitting.

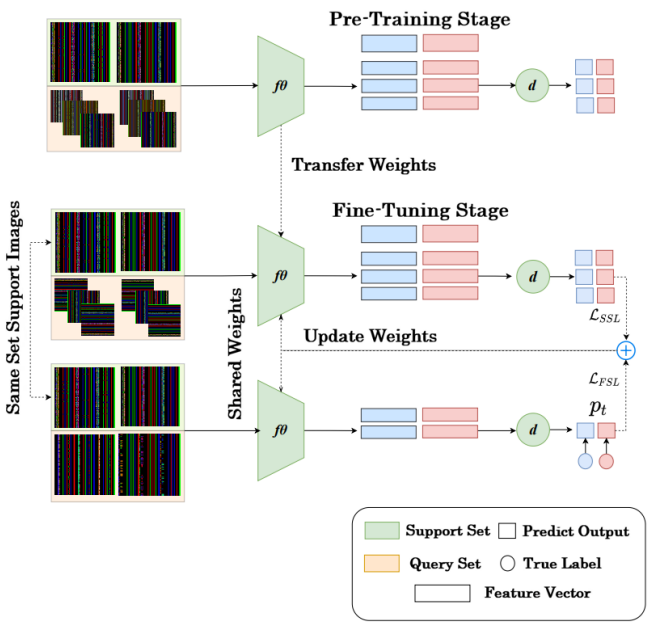

Fig. 3: Architecture of the proposed SSL–FSL model. A shared encoder is used across both pre-training and fine-tuning stages.

Let the training dataset be denoted as $\mathcal{D}$. A few-shot task is defined by a support set:

$$S = \{(x_1, y_1), \ldots, (x_I, y_I)\} \tag{1}$$

where x_i is an input image and $y_i \in \{1, \ldots, N\}$ is its label. The support and query subsets for class k are denoted as S_k and Q_k, respectively.

1) Self-Supervised Pre-Training: During the self-supervised stage, label information is not used. Each sample acts as its own prototype, and augmented samples are generated using rotation transformations as shown in Fig. 4(a):

$$R = \{90°, 180°, 270°\} \tag{2}$$

The rotation set is limited to 90°, 180°, and 270° to ensure effective augmentation without redundancy, as 0° would result in the original sample; therefore, it does not provide an additional learning signal.

For each augmented sample $\hat{x}_{i,q}$, the encoder f_θ minimises the squared Euclidean distance:

$$d(f_\theta(\hat{x}_{i,q}), f_\theta(x_i)) = \|f_\theta(\hat{x}_{i,q}) - f_\theta(x_i)\|_2^2 \tag{3}$$

This distance enforces consistency between original and augmented samples, encouraging the encoder to learn transformation-invariance. A probability distribution is computed via softmax over negative distances, yielding a log-softmax (cross-entropy) formulation:

$$\lambda(i,q) = -\log \frac{\exp\left(-d(f_\theta(\hat{x}_{i,q}), f_\theta(x_i))\right)}{\sum_{k=1}^{M} \exp\left(-d(f_\theta(\hat{x}_{i,q}), f_\theta(x_k))\right)} \tag{4}$$

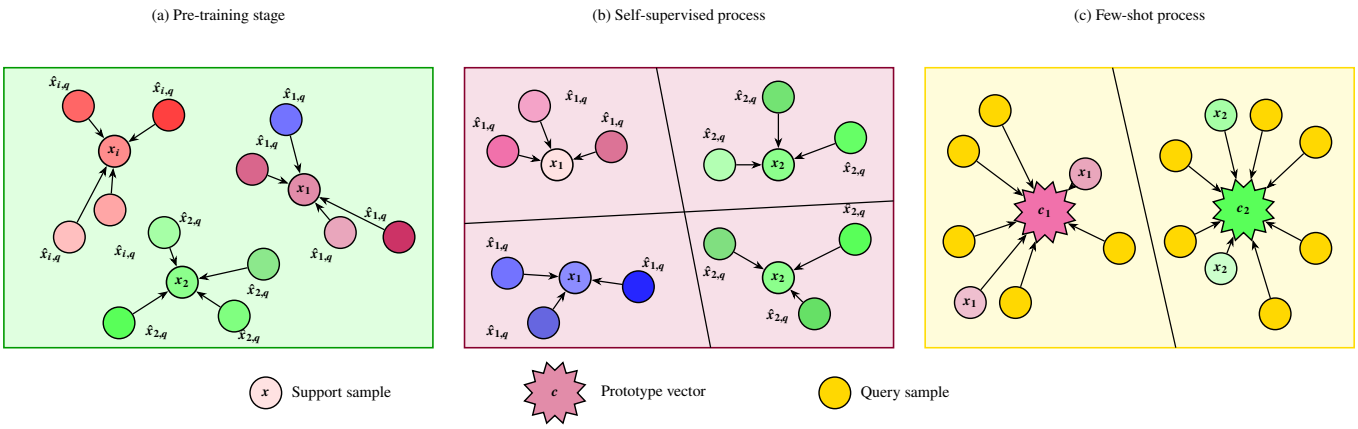

Fig. 4: Conceptual overview of the proposed SSL–FSL framework. (a) self-supervised pre-training, (b) representation refinement, and (c) few-shot classification using class prototypes.

This objective encourages augmented samples to remain close to their original representations while increasing their separation from other samples in the embedding space.

The pre-training loss is defined as:

$$L_{\text{SSL}}^{\text{pre}} = \frac{1}{3M} \sum_{i=1}^{M} \sum_{q=1}^{3} \lambda(i, q) \tag{5}$$

2) Few-Shot Fine-Tuning:

a) Prototype Construction:

$$c_k = \frac{1}{|S_k|} \sum_{(x_i, y_i) \in S_k} f_\theta(x_i) \tag{6}$$

The prototype c_k represents the centroid of class k in the embedding space and serves as a reference for classification.

b) Few-Shot Classification:

$$L_{\text{FSL}} = \frac{1}{|Q_k|} \sum_{x_i \in Q_k} \Big[d(f_\theta(x_i), c_k) + \log \sum_{k'} \exp\left(-d(f_\theta(x_i), c_{k'})\right) \Big] \tag{7}$$

Query samples are classified based on their distance to class prototypes, as shown in Fig. 4(c). Where smaller distances indicate higher similarity and higher classification confidence.

c) Self-Supervised Refinement:

$$L_{\text{SSL}} = \frac{1}{3|S_k|} \sum_{i=1}^{|S_k|} \sum_{q=1}^{3} \lambda(i, q) \tag{8}$$

This auxiliary objective reinforces representation consistency during fine-tuning and helps reduce overfitting under limited labeled data, as illustrated in Fig. 4(b).

d) Final Objective Function:

$$L = \beta L_{\text{SSL}} + L_{\text{FSL}} \tag{9}$$

where $\beta \in [0, 1]$ controls the balance between self-supervised regularization and prototype-based classification. Smaller β values generally yield better performance by prioritizing class discrimination while retaining regularization, whereas larger values may overemphasize consistency and reduce class separability. β acts as a regularization coefficient rather than a convex weight: L_{FSL} is the primary objective (unit weight), while L_{SSL} serves as an auxiliary regularizer. Thus, the objective is defined as $L = \beta L_{\text{SSL}} + L_{\text{FSL}}$ to preserve classification performance while benefiting from representation regularization.

IV. Experimental Setup

The proposed framework is evaluated on the Edge-IIoTset [12] and the unseen CIC-IoT2023 dataset [13]. A subset of the Edge-IIoT dataset is used to train the encoder, and the remaining classes are used during the few-shot stage. The CICIoT2023 dataset is used during the evaluation phase as the model is encountering unseen data. The preprocessing steps include removing duplicate samples (815 samples, < 0.04%), eliminating 16 features that did not contain any informative values for the models, one-hot encoding 7 categorical features, and normalizing each feature. Additionally, all samples are converted into images to be used within the proposed framework.

Experiments are performed on a GPU-enabled high-performance computing environment. The model is trained using the Adam optimizer with a learning rate 1×10^{-4} in an episodic learning fashion with N-way K-shot settings using rotation-based self-supervised training and few-shot fine-tuning as optimized by Equation 9.

A. Evaluation Framework

Evaluation focuses on the generalization capability of the proposed SSL-FSL framework when applied to previously unseen attack categories. The results are presented using confusion matrices and t-SNE visualizations to assess classification performance and the separability of learned feature representations under few-shot settings. In addition, an ablation study on the SSL weighting parameter β is conducted to analyze its impact on model stability and performance.

β	1-shot	5-shot	10-shot
0.00	88.90	89.52	92.32
0.01	93.45	94.11	96.56
0.03	92.50	93.07	95.57
0.05	91.01	91.81	95.44
0.07	90.97	91.56	94.98
0.10	90.16	91.10	94.26

TABLE I: Effect of the SSL weighting parameter β on accuracy (%).

The results show that a small SSL contribution provides the best balance between representation learning and prototype adaptation.

Furthermore, varying the number of shots demonstrates improved performance as more support samples become available. Cross-dataset evaluation on the completely unseen CICIoT2023 dataset further highlights the model ability to generalize under domain shift without fine-tuning. Additionally, the effect of varying the number of ways in the support set is evaluated, showing that performance decreases as task complexity increases.

V. Results and Discussion

This section presents the experimental evaluation of the proposed SSL-FSL framework, focusing on generalization to attack categories under limited labeled data conditions.

A. Few-Shot Learning Performance

The model is trained using self-supervised pre-training followed by episodic few-shot fine-tuning. As shown in Fig. 5, performance consistently improves with the number of shots (e.g., 10 shots), but it also introduces additional computational overhead. Fig. 5 illustrates the separation of test classes in the embedding space and the corresponding classification performance.

To further analyze this behavior, the effect of the self-supervised regularization weight β is evaluated, as presented in Table I.

Conversely, removing SSL regularization entirely ($\beta = 0$) results in a drop in accuracy, confirming its critical role in stabilizing representations for unseen-class generalization. Although the 10-shot setting still outperforms the 1-shot and 5-shot configurations under $\beta = 0$, it remains inferior to the best-performing $\beta = 0.01$ case. Notably, the 10-shot accuracy without SSL (92.32%) is lower than the 1-shot performance with SSL (93.45%), demonstrating that SSL regularization contributes more to representation quality than simply increasing the number of support samples and that its benefits persist even in higher-shot regimes.

B. Generalization to Unseen Dataset

Following the transformation strategy of Latif *et al.* [11], the CICIoT2023 dataset is converted into image representations and reshaped into $46 \times 46 \times 3$ RGB images. The model is evaluated on the completely unseen dataset without fine-tuning under multiple N-way classification settings as presented in Table II.

Task Setting	**Accuracy (%)**
2-way classification	98.12
3-way classification	97.75
5-way classification	92.08
9-way classification	90.49
Full (34-way classification)	69.76

TABLE II: Effect of task complexity on classification accuracy.

In few-shot learning, an N-way classification task refers to a scenario where the model must distinguish between N different classes using a limited number of labeled examples. For instance, 2-way, 3-way, 5-way, and 9-way classification correspond to tasks involving 2, 3, 5, and 9 classes, respectively. The full setting represents classification across all available unseen classes in the dataset.

C. Comparison with Existing Method

To provide a fair comparison, we evaluate the proposed method against an attention-based Siamese network [14], a representative supervised few-shot learning approach widely used in intrusion detection.

As shown in Table III, the proposed SSL-FSL framework achieves higher accuracy than the attention-based Siamese model in all settings. In the 1-shot setting, the SSL-FSL framework achieves 93.45% accuracy compared to 60.22% accuracy in [14]. Furthermore, the model maintains its accuracy with an increase in the number of shots provided to the model. For example, when the number of shots is increased to 10, the accuracy for the proposed framework reaches 96.56% compared to 91.41% for the attention-based Siamese model. Furthermore, the framework is evaluated in more challenging settings (with an increasing number of classes). In these evaluations, the framework maintains its accuracy, indicating its robustness relative to the Siamese model. Thus, these results provide evidence of the effectiveness of the proposed framework relative to an existing few-shot learning framework.

VI. Discussion

The proposed SSL-FSL framework learns robust feature representations through self-supervised pre-training, enabling effective performance even in low-shot settings. The model can also benefit from additional support samples while maintaining its performance with increasing task complexity. Most intrusion detection research relies on publicly available datasets. While evaluating the model on independently generated novel attacks would be ideal, such experiments require significant resources

Model	**1-shot (%)**	**5-shot (%)**	**10-shot (%)**	**Learning Strategy**
Siamese + Attention [14]	60.22	71.82	72.22	FSL (Supervised)
Proposed SSL-FSL	**93.45**	**94.11**	**96.56**	**SSL + FSL**

TABLE III: Comparison on Edge-IIoTset under 5-way Few-Shot Setting

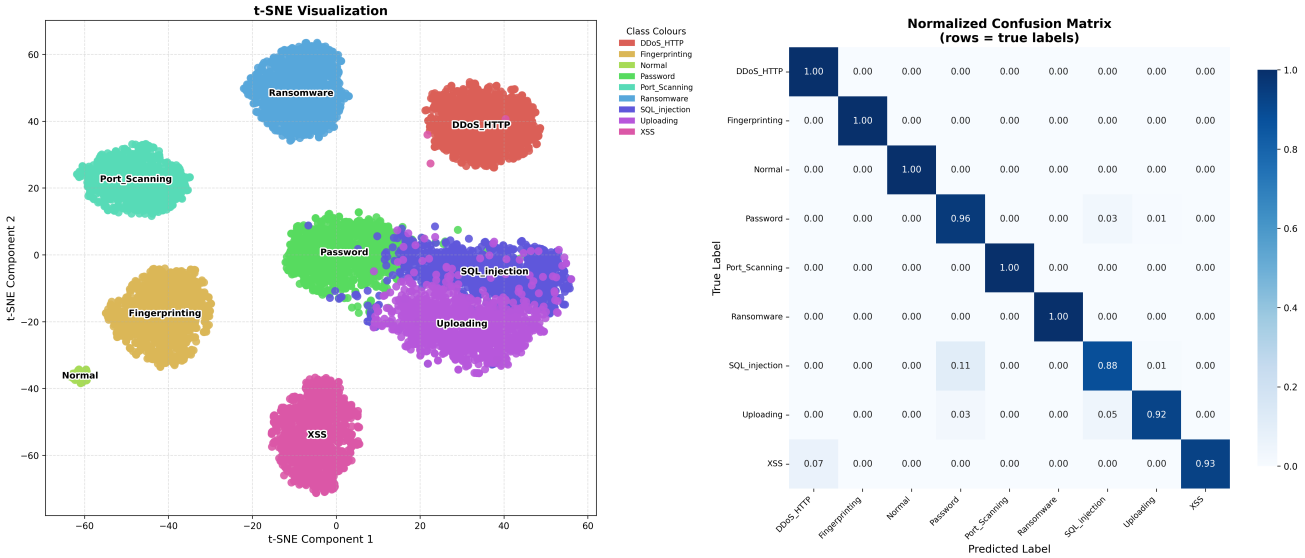

Fig. 5: t-SNE visualization and normalized confusion matrix for test classes under the 10-shot setting using Edge-IIoT.

to perform. To mitigate the impact of relying on publicly available datasets, the proposed model is evaluated on a completely unseen dataset (CICIoT2023). The incorporation of SSL and prototype classification leads to better generalization capabilities due to the learning of features that are more transferable to other network traffic classification tasks.

VII. Conclusion

This paper introduced an SSL-FSL framework for IoT intrusion detection, achieving 93.45% (1-shot), 94.11% (5-shot), and 96.56% (10-shot), while maintaining robust performance under increasing task complexity. The model also attained 69.76% accuracy on completely unseen attacks without fine-tuning on the CICIoT2023 dataset, demonstrating strong cross-dataset generalization. Furthermore, it significantly outperformed baseline methods (e.g., +33% improvement in 1-shot settings) on the Edge-IIoTset dataset, confirming the effectiveness of combining self-supervised and few-shot learning for data-scarce IoT environments.

Acknowledgment

The authors would like to thank King Fahd University of Petroleum and Minerals (KFUPM) for providing the research environment and facilities that made this work possible.

References

[1] A. Thakkar and R. Lohiya, "A review on machine learning and deep learning perspectives of ids for iot: Recent updates, security issues, and challenges." *Archives of computational methods in engineering*, vol. 28, no. 4, 2021.

[2] H. Liu and B. Lang, "Machine learning and deep learning methods for intrusion detection systems: A survey," *applied sciences*, vol. 9, no. 20, p. 4396, 2019.

[3] A. Abualhassan, I. Rashid, F. Binbeshr, and M. Imam, "DDos attack detection in iot: A comparative resource and performance analysis of deep learning and machine learning models," *IEEE Access*, vol. 13, pp. 116 529–116 547, 2025.

[4] A. Althnian, D. AlSaeed, H. Al-Baity, A. Samha, A. B. Dris, N. Alzakari, A. Abou Elwafa, and H. Kurdi, "Impact of dataset size on classification performance: An empirical evaluation in the medical domain," *Applied Sciences*, vol. 11, no. 2, 2021. [Online]. Available: https://www.mdpi.com/2076-3417/11/2/796

[5] A. A.Abualhassan, Y. Fadol, M. S. M. Gismalla, and M. Hamdan, "Iiot-tinydnn: A lightweight intrusion detection system for edge-based iiot security," in *2025 IEEE Conference on Standards for Communications and Networking (CSCN)*, 2025, pp. 1–6.

[6] A. Dina, A. Siddique, and D. Manivannan, "Fs3: Few-shot and self-supervised framework for efficient intrusion detection in internet of things networks," 12 2023, pp. 138–149.

[7] S. B. Atitallah, M. Driss, W. Boulila, and A. Koubaa, "Strengthening network intrusion detection in iot environments with self-supervised learning and few shot learning," 2024. [Online]. Available: https://arxiv.org/abs/2406.02636

[8] C. Xu, Y. Zhan, G. Chen, Z. Wang, S. Liu, and W. Hu, "Elevated few-shot network intrusion detection via self-attention mechanisms and iterative refinement," *PLOS ONE*, vol. 20, no. 1, pp. 1–22, 01 2025. [Online]. Available: https://doi.org/10.1371/journal.pone.0317713

[9] S. Gidaris, A. Bursuc, N. Komodakis, P. PÃ©rez, and M. Cord, "Boosting few-shot visual learning with self-supervision," 2019. [Online]. Available: https://arxiv.org/abs/1906.05186

[10] J. Y. Lim, K. M. Lim, C. P. Lee, and Y. X. Tan, "Ssl-protonet: Self-supervised learning prototypical networks for few-shot learning," *Expert Systems with Applications*, vol. 238, p. 122173, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417423026751

[11] S. Latif, W. Boulila, A. Koubaa, Z. Zou, and J. Ahmad, "Dtl-ids: An optimized intrusion detection framework using deep transfer learning and genetic algorithm," *Journal of Network and Computer Applications*, vol. 221, p. 103784, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1084804523002035

[12] M. A. Ferrag, O. Friha, D. Hamouda, L. Maglaras, and H. Janicke, "Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning," *IEEE Access*, vol. 10, pp. 40 281–40 306, 2022.

[13] E. C. P. Neto, S. Dadkhah, R. Ferreira, A. Zohourian, R. Lu, and A. A. Ghorbani, "Ciciot2023: A real-time dataset and benchmark for large-scale attacks in iot environment," *Sensors*, vol. 23, no. 13, 2023. [Online]. Available: https://www.mdpi.com/1424-8220/23/13/5941

[14] S. Latif, J. Ahmad, W. Boulila, M. S. Khan, and D. Djenouri, "Few-shot learning for iot intrusion detection: An attention-based siamese network approach," *Procedia Computer Science*, vol. 270, pp. 4553–4562, 2025, 29th International Conference on Knowledge-Based and Intelligent Information Engineering Systems (KES 2025). [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050925032545