

MIETT: Multi-Instance Encrypted Traffic Transformer for Encrypted Traffic Classification

Xu-Yang Chen*, Lu Han*, De-Chuan Zhan, Han-Jia Ye[†]

School of Artificial Intelligence, Nanjing University, China
National Key Laboratory for Novel Software Technology, Nanjing University, China
{chenxy, hanlu, zhandc, yehj}@lamda.nju.edu.cn

Abstract

Network traffic includes data transmitted across a network, such as web browsing and file transfers, and is organized into packets (small units of data) and flows (sequences of packets exchanged between two endpoints). Classifying encrypted traffic is essential for detecting security threats and optimizing network management. Recent advancements have highlighted the superiority of foundation models in this task, particularly for their ability to leverage large amounts of unlabeled data and demonstrate strong generalization to unseen data. However, existing methods that focus on token-level relationships fail to capture broader flow patterns, as tokens, defined as sequences of hexadecimal digits, typically carry limited semantic information in encrypted traffic. These flow patterns, which are crucial for traffic classification, arise from the interactions between packets within a flow, not just their internal structure. To address this limitation, we propose a Multi-Instance Encrypted Traffic Transformer (MIETT), which adopts a multi-instance approach where each packet is treated as a distinct instance within a larger bag representing the entire flow. This enables the model to capture both token-level and packet-level relationships more effectively through Two-Level Attention (TLA) layers, improving the model’s ability to learn complex packet dynamics and flow patterns. We further enhance the model’s understanding of temporal and flow-specific dynamics by introducing two novel pre-training tasks: Packet Relative Position Prediction (PRPP) and Flow Contrastive Learning (FCL). After fine-tuning, MIETT achieves state-of-the-art (SOTA) performance across five datasets, demonstrating its effectiveness in classifying encrypted traffic and understanding complex network behaviors.

Introduction

Network traffic refers to the flow of data transmitted between devices over a network, typically structured into packets, which are small units sent across the network, and flows, which are sequences of packets exchanged between two points. Each packet is composed of two parts: the header and the payload. The header contains essential information, such as routing details, source and destination addresses, and

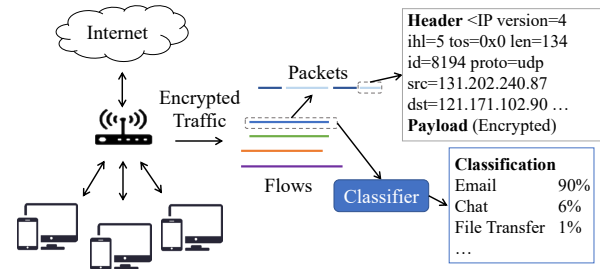


Figure 1: Encrypted traffic classification task description. Raw traffic is first divided into session flows, with each flow further segmented into a sequence of packets. A packet typically consists of a header and a payload. The task is to classify the type of a given flow.

packet length, while the payload carries the actual data being transmitted, which may be encrypted for security purposes.

Traffic classification, the process of identifying and categorizing network traffic, is crucial for both network management and cybersecurity. It allows network administrators to ensure Quality of Service (QoS), optimize bandwidth and detect malicious activities, thus maintaining network security and efficiency. The overview of encrypted traffic classification task is provided in Figure 1.

However, the increasing prevalence of encryption has made traditional classification methods, such as port-based and statistics-based approaches, less effective. The advent of deep learning (DL) brought significant improvements, with payload-based methods using models like CNNs to automatically extract features from raw data. Despite their success, these methods rely heavily on large amounts of labeled data, which can be difficult to obtain.

Recently, foundation models have emerged as a powerful alternative, pre-trained on large volumes of unlabeled data and fine-tuned for specific tasks. PERT (He, Yang, and Chen 2020) employs a BERT-based model with a Masked Language Modeling (MLM) task to pre-train a packet-level encoder, but this approach primarily focuses on token-level relationships within individual packets, overlooking the broader context of inter-packet relationships. ET-BERT (Lin et al. 2022) addresses this limitation by introducing the Same-origin BURST Prediction (SBP) task, which deter-

*These authors contributed equally.

[†]Corresponding Author

mines whether one packet follows another within the same flow. While this method accounts for adjacent packet relationships, it still falls short of capturing the full complexity of flow-level interactions and the broader context across multiple packet segments. YaTC (Zhao et al. 2023) takes a different approach by tokenizing traffic data into patches and using a Masked Auto-Encoder (MAE) for pre-training a token-level encoder within a flow. However, focusing solely on token dependencies is less effective due to the low semantic information in encrypted traffic tokens. This approach tends to overlook the broader patterns and relationships that exist between entire packets within a traffic flow. Instead, learning packet representations that capture these packet patterns proves to be more robust. Given that each packet can be viewed as a distinct instance carrying unique information within a flow, it is crucial to effectively model the relationships between packets to achieve a comprehensive flow representation. To address this, we propose a Multi-Instance Encrypted Traffic Transformer (MIETT) with Two-Level Attention (TLA) layers that capture both token-level and packet-level relationships. To further enhance the model’s ability to capture temporal and flow-specific dynamics, we introduce two novel pre-training tasks: Packet Relative Position Prediction (PRPP) and Flow Contrastive Learning (FCL). The PRPP task is designed to help the model understand the sequential relationships between packets within a flow by predicting their relative positions, thereby enabling a more accurate representation of the flow’s structure. Meanwhile, the FCL task focuses on distinguishing packets from the same flow and those from different flows by learning robust representations that emphasize intra-flow similarities and inter-flow differences.

Building on the robust representations learned during pre-training, the model is then fully fine-tuned on the specific classification task. During this stage, both the packet encoder and the flow encoder are jointly optimized to adapt the model to the task at hand. Across experiments conducted on five datasets, MIETT consistently demonstrates competitive or superior performance compared to existing methods. In conclusion, this paper presents several key contributions to the field of encrypted traffic classification:

- We propose a novel Multi-Instance Encrypted Traffic Transformer (MIETT) architecture, which introduces Two-Level Attention (TLA) layers to effectively capture both token-level and packet-level relationships within traffic flows.
- We introduce two innovative pre-training tasks, Packet Relative Position Prediction (PRPP) and Flow Contrastive Learning (FCL). The PRPP task enhances the model’s understanding of the sequential order of packets within a flow, while the FCL task improves the model’s ability to differentiate between packets from the same flow and those from different flows.
- We provide extensive empirical validation of the proposed MIETT model on multiple datasets, showing that our approach outperforms existing methods in terms of accuracy and F1-score.

Related Work

Encrypted traffic classification is an essential research area in network security. As encryption becomes more prevalent, traditional traffic analysis methods that inspect packet payloads have become less effective. Researchers now focus on methods to classify encrypted traffic without accessing its content directly. These methods can be divided into three types: port-based methods, statistics-based methods, and payload-based methods.

Port-based Methods. Port-based traffic classification, one of the oldest methods, relies on associating port numbers in TCP/UDP headers with well-known IANA port numbers. Although it is fast and simple, the method’s accuracy has declined due to port obfuscation, dynamic port assignments, network address translation (NAT), and other techniques (Moore and Papagiannaki 2005).

Statistics-based Methods. Statistics-based methods for encrypted traffic classification leverage features independent of payloads, such as packet sizes, timing, and flow duration, to analyze and categorize traffic. Wang et al. computed entropy from packet payloads to classify eight traffic classes, using an SVM algorithm to select features (Wang et al. 2011). Similarly, Korczynski and Duda focused on packet sizes, timing, and communication patterns for classifying traffic in encrypted tunnels (Korczynski and Duda 2012). However, these methods are based on manually designed features and are now outdated due to advancements in encryption protocols, changes in traffic patterns, and the use of traffic obfuscation techniques.

Payload-based DL Methods. Payload-based methods usually leverage deep learning to analyze raw packet data, eliminating the need for manually designed features. These approaches have significantly advanced encrypted traffic classification by automatically extracting discriminative features. Deep Packet (Lotfollahi et al. 2020) employs a CNN within a framework that includes a stacked autoencoder and convolutional neural network to classify network traffic. TSCRNN (Lin, Xu, and Gao 2021) utilizes CNN to extract abstract spatial features and introduces a stacked bidirectional LSTM to learn temporal characteristics. BiLSTM-ATTN (Yao et al. 2019) combines attention mechanisms with LSTM networks for enhanced encrypted traffic classification. Though these methods can automatically extract features, they heavily rely on labeled data, which is often difficult to obtain in large quantities for training.

Payload-based Foundation Models. In recent years, pre-trained foundation models have become popular for addressing this issue. These models are pre-trained on large amounts of unlabeled data and fine-tuned on downstream labeled tasks. PERT (He, Yang, and Chen 2020) and ET-BERT (Lin et al. 2022) tokenize traffic data using a vocabulary. PERT employs a Masked Language Model (MLM) pre-training task, while ET-BERT utilizes a modified MLM and Next Sentence Prediction (NSP) task. Additionally, YaTC (Zhao et al. 2023), a vision model adaptation, tokenizes traffic data into patches and uses a Mask Auto-Encoder (MAE) for pre-training. However, these approaches do not adequately con-

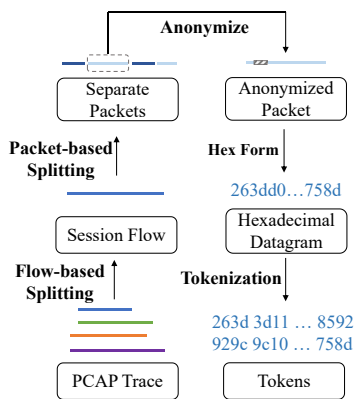


Figure 2: Data preprocessing. The raw traffic (PCAP trace) is first split into session flows and then further divided into individual packets. To protect data privacy, each packet is anonymized by masking the source and destination IP addresses and port numbers (replacing them with 0). The packet is then converted to its hexadecimal form, which is tokenized using a bi-gram model.

consider the unique structure of traffic flows and the relationships between packets. To address these limitations, we introduce the Multi-Instance Encrypted Traffic Transformer (MIETT) architecture, which leverages novel Packet Relative Position Prediction (PRPP) and Flow Contrastive Learning (FCL) tasks to better capture the complexities of traffic flows.

Multi-Instance Encrypted Traffic Transformer

In the task of encrypted traffic classification, we are provided with raw network traffic (PCAP traces) as input, and the objective is to classify it into categories, such as VPN services (e.g., P2P, streaming, email) and applications. In this section, we first outline the preprocessing steps that transform the raw data into a multi-instance traffic representation. We then introduce our Multi-Instance Encrypted Traffic Transformer (MIETT) architecture, which is specifically designed to handle and classify this traffic data efficiently.

MIETT Encoder

This section details the process of representing multi-instance traffic data for use in the Multi-Instance Encrypted Traffic Transformer (MIETT) architecture. The process involves three key steps: tokenization of the raw data, representation of individual packets, and the aggregation of these packet representations into a unified flow representation.

Tokenization. The hexadecimal sequence of the flow is obtained through the data preprocessing steps outlined in Figure 2. Following ET-BERT (Lin et al. 2022), we encode the hexadecimal sequence using a bi-gram model, where each unit consists of two consecutive bytes. We then utilize Byte-Pair Encoding (BPE) for token representation, with token units ranging from 0 to 65535 and a maximum dictionary size of 65536. For training, we also incorporate special tokens, including [CLS], [PAD], and [MASK].

Packet Representation. Each packet begins with a [CLS] token, followed by tokens extracted from the packet’s content, which includes a header containing meta-features and an encrypted payload. The embedding of each token is obtained by combining two parts: position embedding, which indicates the token’s position within the packet, and value embedding. To ensure efficient information utilization and avoid excessive focus on long packets, the packet length is standardized to a fixed size of 128. Packets shorter than 128 tokens are padded with [PAD] tokens at the end.

Flow Representation. The flow representation consists of multiple packet representations, which are stacked to form a matrix $\mathbf{X} \in \mathbb{R}^{N \times L \times d}$, where N is the number of packets, L is the packet length, and d is the embedding dimension. This multi-instance representation, as opposed to the previous method employed by ET-BERT of directly concatenating packets, allows for more effective modeling of the relationships between packets and better captures the organizational structure of flows.

MIETT Architecture

The flow representation, a 2D map of tokens, can be flattened into a 1D sequence for input into a standard transformer, as done in ViT (Dosovitskiy et al. 2021) and YaTC (Zhao et al. 2023) during pre-training and fine-tuning. However, this approach introduces two issues: (1) Flattening the sequence loses temporal information, such as packet order, potentially overlooking important dependencies. While 2D position embeddings could indicate temporal relationships, they are inflexible and may fail in scenarios like packet loss. (2) The computational complexity, which is $O(N^2L^2d)$, increases significantly with the inclusion of more packets due to the extended sequence length. To address these concerns, we introduce the Two-Level Attention (TLA) layer to preserve temporal structure and maintain computational efficiency.

Overall Architecture. The Multi-Instance Encrypted Traffic Transformer (MIETT) begins by embedding the tokens as described in the MIETT Encoder section. Following this, the flow representation is processed through M TLA layers. Finally, the embeddings of the [CLS] tokens from all the packets are utilized for pre-training or fine-tuning tasks. The overall architecture of MIETT can be found in Figure 3.

Two-Level Attention (TLA) Layer. The TLA layer captures both intra-packet and inter-packet dependencies to enhance the model’s understanding of complex traffic flows. It operates in two stages: Packet Attention and Flow Attention.

In the Packet Attention stage, Multi-Head Self-Attention (MHSA) is applied within individual packets to identify dependencies between tokens, ensuring the model understands the internal structure of each packet. For MHSA, the key, query, and value are all set to the input sequence, with the output being an enriched sequence representation.

In the Flow Attention stage, MHSA is used across packet representations at each position to capture dependencies between tokens across different packets. This two-stage approach effectively models the hierarchical structure of traf-

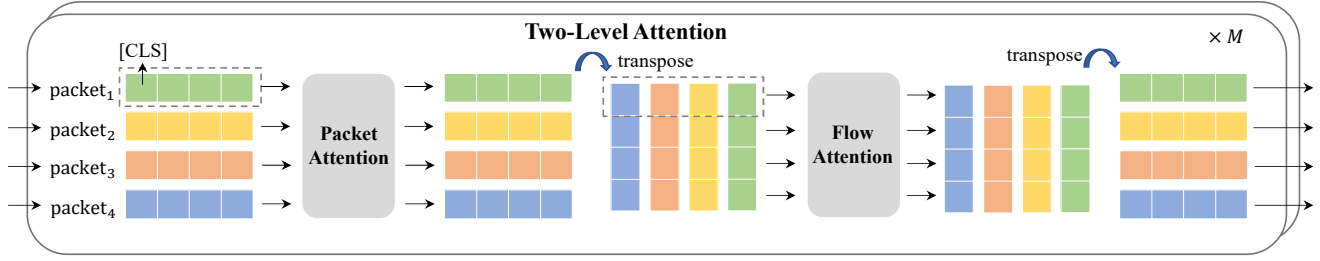


Figure 3: Overall architecture of the Multi-Instance Encrypted Traffic Transformer (MIETT). After passing through the MIETT encoder, the flow representation is processed by M Two-Level Attention (TLA) layers, each comprising a packet attention mechanism and a flow attention mechanism.

fic flows by combining detailed packet-level insights with broader inter-packet relationships.

This method significantly improves the model’s ability to capture complex dependencies but also introduces computational complexity. The MHSA in Packet Attention has a complexity of $O(L^2d)$ per packet, while Flow Attention has a complexity of $O(N^2d)$ per position, resulting in an overall complexity of $O(NL^2d + LN^2d)$. For default values of $L = 128$ and $N = 5$, our method is approximately 4.8 times more efficient than flattening the sequence to 1D and feeding it into a standard transformer.

Packet Attention. In the first stage of the TLA layer, we focus on the intra-packet relationships by performing Multi-Head Self-Attention (MHSA) within individual packets. Let $\mathbf{X} \in \mathbb{R}^{N \times L \times d}$ be the flow representation, where N is the number of packets, L is the packet length, and d is the embedding dimension.

For each packet $\mathbf{X}_i \in \mathbb{R}^{L \times d}$, we apply self-attention to capture the dependencies between the tokens within the packet. The process is as follows:

$$\hat{\mathbf{X}}_i^{\text{pkt}} = \text{LayerNorm}(\mathbf{X}_i + \text{MHSA}^{\text{pkt}}(\mathbf{X}_i)) \quad (1)$$

$$\mathbf{X}_i^{\text{pkt}} = \text{LayerNorm}(\hat{\mathbf{X}}_i^{\text{pkt}} + \text{MLP}(\hat{\mathbf{X}}_i^{\text{pkt}})) \quad (2)$$

Flow Attention. In the second stage of the TLA layer, we focus on the inter-packet relationships by performing multi-head self-attention (MHSA) across the packet representations at each position within the packets. Let $\mathbf{X}^{\text{packet}} \in \mathbb{R}^{N \times L \times d}$ be the updated flow representation after the packet attention stage.

For each position j (where $j \in 1, 2, \dots, L$) within the packets, we gather the token representations across all packets, resulting in a matrix $\mathbf{X}_j^{\text{packet}} \in \mathbb{R}^{N \times d}$, where N is the number of packets. We apply self-attention to these matrices to capture the dependencies between tokens across different packets. The process is as follows:

$$\hat{\mathbf{X}}_j^{\text{flow}} = \text{LayerNorm}(\mathbf{X}_j^{\text{pkt}} + \text{MHSA}^{\text{flow}}(\mathbf{X}_j^{\text{pkt}})) \quad (3)$$

$$\mathbf{X}_j^{\text{flow}} = \text{LayerNorm}(\hat{\mathbf{X}}_j^{\text{flow}} + \text{MLP}(\hat{\mathbf{X}}_j^{\text{flow}})) \quad (4)$$

Training Tasks

This section outlines the training tasks designed to improve the model’s ability to classify encrypted network traffic. The

pre-training phase includes three tasks: Masked Flow Prediction (MFP), Packet Relative Position Prediction (PRPP), and Flow Contrastive Learning (FCL). These tasks help the model capture flow dependencies, predict packet order, and differentiate flow-level features. After pre-training, the model is fine-tuned for traffic flow classification, optimizing it for the final task.

Pre-Training Tasks

The model we propose consists of several Two-Level Attention (TLA) layers, designed to effectively capture both packet-level and flow-level information within traffic flows. During the pre-training stage, We utilize a pre-trained E-BERT checkpoint for the packet attention, which is kept frozen during training, while the flow attention is trained to learn the overall structure and dependencies within the flows. This design allows our model to leverage established packet-level features while focusing on improving flow-level understanding. The general view of our proposed 3 pre-training tasks can be found in Figure 4.

Masked Flow Prediction (MFP) Task. The Masked Flow Prediction (MFP) task is aimed at enhancing the model’s ability to handle incomplete information within a traffic flow. In this task, 15% of the tokens within a flow are randomly masked, and the model is tasked with predicting the original content of these masked tokens using the context provided by the unmasked tokens. By training the flow encoder to infer the missing tokens, the model learns to capture the underlying structure and dependencies within the flow.

Packet Relative Position Prediction (PRPP) Task. The Packet Relative Position Prediction (PRPP) task is designed to predict the relative order of packets within a flow, based on the embeddings of the [CLS] tokens extracted from each packet. This task helps the model understand the sequential relationships between packets, which is critical for accurately modeling traffic flows.

Given an output flow representation with N packets, let $\mathbf{O}^{\text{pkt}} \in \mathbb{R}^{N \times d}$ represent the embeddings of the [CLS] tokens from all packets in the flow, where d is the embedding dimension. The task is to determine, for each pair of packets (i, j) , whether packet i comes before packet j .

Firstly, a linear transformation is applied followed by an

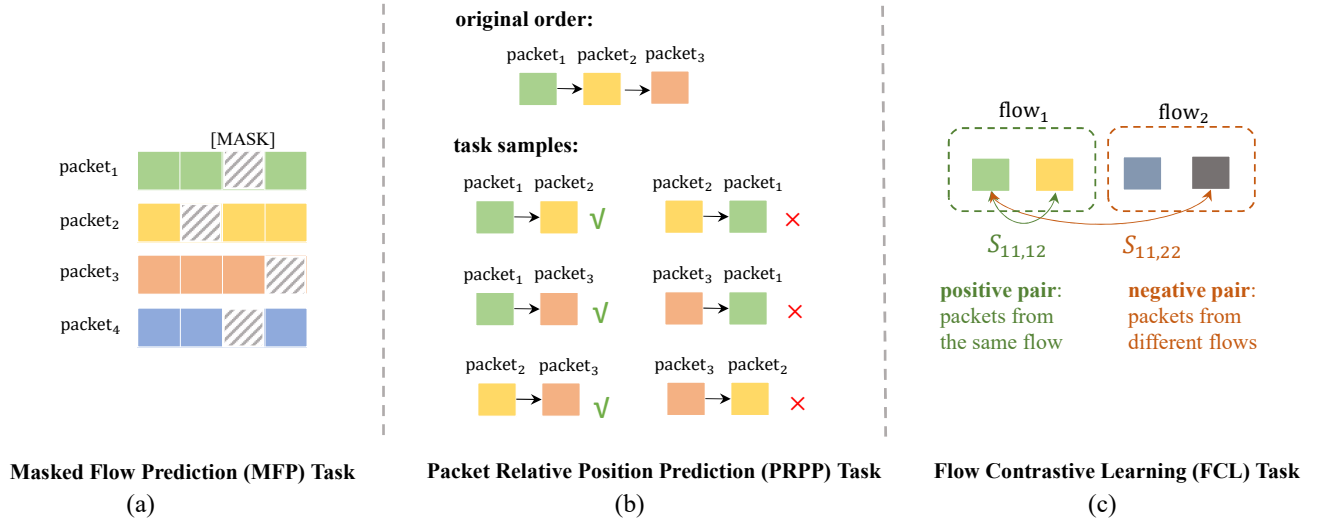


Figure 4: Overview of pre-training tasks. (a) Masked Flow Prediction (MFP) Task: The model is tasked with predicting the original content of masked tokens using the context provided by the unmasked tokens. (b) Packet Relative Position Prediction (PRPP) Task: The model’s objective is to determine, for each pair of packets (i, j), whether packet i precedes packet j . (c) Flow Contrastive Learning (FCL) Task: The goal is to ensure that packets within the same flow (positive pairs) are more similar in the embedding space, while packets from different flows (negative pairs) are less similar.

activation function and layer normalization to [CLS] tokens:

$$\mathbf{P} = \text{LayerNorm}(\text{GELU}(\mathbf{O}^{\text{pkt}}\mathbf{W}_1 + \mathbf{b}_1)) \quad (5)$$

where $\mathbf{W}_1 \in \mathbb{R}^{d \times d}$ and $\mathbf{b}_1 \in \mathbb{R}^d$ are learnable parameters. $\mathbf{P} \in \mathbb{R}^{N \times d}$.

Then, the predicted relative position of packets (i, j) is computed as follows:

$$\hat{z}_{ij} = \text{Softmax}((\mathbf{P}_i - \mathbf{P}_j)\mathbf{W}_2 + \mathbf{b}_2) \quad (6)$$

where $\mathbf{W}_2 \in \mathbb{R}^{d \times 2}$ and $\mathbf{b}_2 \in \mathbb{R}^2$ are learnable parameters, and $\hat{z}_{ij} \in \mathbb{R}^2$ gives the probability that packet i comes before or after packet j .

The ground truth labels $z_{ij} \in \{0, 1\}$ are based on the original order of packets:

$$z_{ij} = \begin{cases} 1 & \text{if packet } i \text{ comes before packet } j \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Finally, the PRPP loss $\mathcal{L}_{\text{PRPP}}$ is calculated using cross-entropy between the predicted labels and the ground truth:

$$\mathcal{L}_{\text{PRPP}} = - \sum_{i,j,i \neq j} z_{ij} \log(\hat{z}_{ij}) + (1 - z_{ij}) \log(1 - \hat{z}_{ij}) \quad (8)$$

Flow Contrastive Learning (FCL) Task. The Flow Contrastive Learning (FCL) task enhances the model’s ability to differentiate between traffic flows by learning robust representations. The objective is to ensure that packets within the same flow (positive pairs) are more similar in the embedding space, while packets from different flows (negative pairs) are less similar. Notably, both positive and negative pairs are constructed using identical packet positions within their respective flows, maintaining consistency in the comparisons.

Given a batch of output flow representations, let $\mathbf{O}^{\text{flow}} \in \mathbb{R}^{BS \times N \times d}$ represent the embeddings of the [CLS] tokens from all packets in the flows within the batch, where BS denotes the batch size, N the number of packets per flow, and d the embedding dimension.

First, a multi-layer perceptron (MLP) is applied to each [CLS] token:

$$\mathbf{C} = \text{LayerNorm}(\text{GELU}(\mathbf{O}^{\text{flow}}\mathbf{W}_3 + \mathbf{b}_3)) \quad (9)$$

$$\mathbf{C} = \mathbf{C}\mathbf{W}_4 + \mathbf{b}_4 \quad (10)$$

where $\mathbf{W}_3 \in \mathbb{R}^{d \times d}$, $\mathbf{b}_3 \in \mathbb{R}^d$, $\mathbf{W}_4 \in \mathbb{R}^{d \times d}$, and $\mathbf{b}_4 \in \mathbb{R}^d$ are learnable parameters. $\mathbf{C} \in \mathbb{R}^{BS \times N \times d}$.

Next, the similarity between two packets is computed using cosine similarity:

$$\mathbf{S}_{i_1 j_1, i_2 j_2} = \frac{\mathbf{C}_{i_1 j_1}^T \mathbf{C}_{i_2 j_2}}{\|\mathbf{C}_{i_1 j_1}\| \|\mathbf{C}_{i_2 j_2}\|} \quad (11)$$

where i_1, i_2 represent the flow IDs within the batch, and j_1, j_2 denote the packet positions within the flow. $\mathbf{C}_{ij} \in \mathbb{R}^d$.

Finally, the contrastive loss is computed using the similarity matrix \mathbf{S} :

$$\mathcal{L}_{\text{FCL}} = - \sum_{\substack{i_1, j_1, j_2 \\ j_1 \neq j_2}} \log \frac{\exp(\mathbf{S}_{i_1 j_1, i_1 j_2})}{\exp(\mathbf{S}_{i_1 j_1, i_1 j_2}) + \sum_{i_2 \neq i_1} \exp(\mathbf{S}_{i_1 j_1, i_2 j_2})} \quad (12)$$

where $i_1, i_2 \in [1, BS]$ represent the flow IDs in the batch, $j_1, j_2 \in [1, N]$ denote the packet positions within the flow.

Conclusion. Overall, the final loss during the pre-train stage is the weighted sum of the above 3 losses:

$$\mathcal{L}_{\text{pt}} = \mathcal{L}_{\text{MPF}} + \alpha \mathcal{L}_{\text{PRPP}} + \beta \mathcal{L}_{\text{FCL}} \quad (13)$$

where α and β are hyperparameters.

Methods	ISCXVPN 2016		ISCXTor 2016		CrossPlatform (Android)		CrossPlatform (iOS)		CICIoT 2023	
	AC	F1	AC	F1	AC	F1	AC	F1	AC	F1
Datanet	69.18%	13.63%	49.81%	9.50%	9.45%	1.53%	4.81%	0.05%	2.50%	0.81%
Fs-Net	29.30%	33.67%	82.03%	63.54%	7.08%	4.11%	10.94%	6.38%	66.80%	54.81%
BiLSTM_ATTEN	0.57%	3.13%	88.33%	55.54%	0.45%	0.04%	0.29%	0.01%	5.79%	4.66%
DeepPacket	69.18%	13.63%	49.81%	9.50%	4.84%	0.04%	4.81%	0.05%	34.35%	8.52%
TSCRNN	69.18%	13.63%	44.70%	13.40%	2.43%	0.24%	2.84%	0.51%	2.50%	0.81%
YaTC	78.05%	70.83%	97.39%	85.12%	<u>91.61%</u>	<u>82.28%</u>	75.31%	69.57%	86.18%	73.16%
ET-BERT	74.62%	<u>71.10%</u>	95.71%	80.29%	<u>84.63%</u>	<u>67.70%</u>	<u>77.05%</u>	<u>74.26%</u>	<u>88.09%</u>	83.29%
MIETT (ours)	<u>76.07%</u>	77.86%	<u>96.60%</u>	<u>82.15%</u>	93.00%	82.36%	79.63%	75.03%	88.53%	<u>82.48%</u>

Table 1: Performance comparisons on encrypted traffic classification tasks. We denote the **best** and second-best results with bold and underline.

Dataset	#Flow	Task	#Label
ISCXVPN 2016	311,390	VPN Service	6
ISCXTor 2016	55,523	Tor Service	7
CrossPlatform (Android)	66,346	Application	212
Cross Platform (iOS)	34,912	Application	196
CICIoT Dataset 2023	1,163,495	IoT Attack	7

Table 2: The statistical information of 5 different datasets.

Fine-Tuning Task

The objective of the fine-tuning task is to classify a given flow into a specific class. After processing the flow through several TLA layers, we obtain a flow representation by extracting the embeddings of the [CLS] tokens from all packets. These embeddings, representing each packet, are then aggregated using mean pooling to form a comprehensive representation of the entire flow. This mean-pooled flow representation is passed through a multi-layer perceptron (MLP) to produce the final classification output, indicating the predicted class of the traffic flow.

During this stage, the entire model, including both the packet encoder (previously frozen during pre-training) and the flow encoder, is fine-tuned. The fine-tuning process optimizes the model by minimizing the cross-entropy loss:

$$\mathcal{L}_{\text{fit}} = - \sum_c y_c \log(\hat{y}_c) \quad (14)$$

where y_c is the true label, and \hat{y}_c is the predicted probability for class c .

Experiments

Experiment Setup

Datasets and Benchmarks. For the encrypted traffic classification task, we evaluate our method on five datasets: ISXVPN 2016 (Draper-Gil et al. 2016), ISXCTor 2016 (Lashkari et al. 2017), and the Cross-Platform (Van Ede et al. 2020) dataset, which includes two

subsets (Android and iOS), as well as the CIC IoT Dataset 2023 (Neto et al. 2023). We utilize data preprocessed by Netbench (Qian et al. 2024). The data used for pre-training consists of the training sets from all five datasets in Netbench, without labels. The dataset statistics and descriptions of the fine-tuning tasks are detailed in Table 2. The data is split into training, validation, and test sets with a ratio of 8:1:1.

Compared Methods. We compare our method against 7 payload-based approaches, including deep learning methods such as Datanet (Wang et al. 2018), Fs-Net (Liu et al. 2019), BiLSTM_ATTEN (Yao et al. 2019), DeepPacket (Lotfollahi et al. 2020), TSCRNN (Lin, Xu, and Gao 2021), as well as foundation models like ET-BERT (Lin et al. 2022) and YaTC (Zhao et al. 2023).

Implementation Details. During the pre-training stage, we set the training steps to 150,000 and randomly select five of the first ten packets for training. The masking ratio for the Masked Flow Prediction (MFP) task is set to 15%. The weights for the Packet Relative Position Prediction (PRPP) and MFP tasks are both set to 0.2. In the fine-tuning stage, we train for 30 epochs using the first five packets. For both stages, the packet length (L) is set to 128, the number of packets (N) is set to 5, the embedding dimension (d) is set to 768, and the number of Two-Level Attention (TLA) layers is set to 12. The learning rate is set to 2×10^{-5} , and the AdamW optimizer is used. All experiments are conducted on a server with two NVIDIA RTX A6000 GPUs.

Main Results

The primary metrics for comparison are Accuracy and F1-Score. Accuracy measures the proportion of correct predictions, while F1-Score balances precision and recall. Table 1 presents the performance comparison of various models on encrypted traffic classification tasks, with baseline results sourced from NetBench (Qian et al. 2024). The results clearly indicate that traditional deep learning methods, such as DataNet, DeepPacket, FS-Net, TSCRNN, and BiLSTM_ATTEN, struggle to generalize effectively to new tasks, particularly on complex datasets. These models frequently exhibit a bias toward dominant classes, resulting in consis-

Methods	CrossPlatform (Android)		CrossPlatform (iOS)	
	AC	F1	AC	F1
from scratch	88.08%	73.62%	71.63%	63.43%
w/o PRPP	90.79%	79.02%	78.96%	74.35%
w/o FCL	91.90%	81.60%	79.46%	74.80%
Ours	93.00%	82.36%	79.63%	75.03%

Table 3: Impact of Pre-Training Tasks.

Methods	CrossPlatform (Android)		CrossPlatform (iOS)	
	AC	F1	AC	F1
w/o pkt attn	62.19%	28.59%	55.58%	39.93%
w/o flow attn	91.85%	80.77%	79.11%	72.46%
TLA (ours)	93.00%	82.36%	79.63%	75.03%

Table 4: Impact of TLA components, where 'pkt attn' refers to packet attention and 'flow attn' refers to flow attention.

tently low F1-scores, especially on datasets like CrossPlatform (Android) and CrossPlatform (iOS).

In contrast, our MIETT model demonstrates significant improvements in both accuracy and F1-scores across the board, showcasing its superior capability to handle the complexities of encrypted traffic. Notably, MIETT consistently achieves competitive or superior performance compared to existing methods. For instance, in the CrossPlatform (Android) dataset, MIETT outperforms ET-BERT with an 8.27% increase in accuracy and a 14.66% increase in F1-score, highlighting the effectiveness of the flow attention.

Ablation Study

Impact of Pre-Training Tasks. Table 3 presents an ablation study that compares different versions of the pre-training tasks on the CrossPlatform (Android) and CrossPlatform (iOS) datasets to evaluate the contributions of specific tasks. The baseline model ("from scratch") shows decent performance, indicating that pre-training is important and can significantly enhance the model's ability to generalize to new tasks. As the results demonstrate, incorporating specific pre-training tasks such as Packet Relative Position Prediction (PRPP) and Flow Contrastive Learning (FCL) further boosts the model's performance.

Impact of TLA Components. TLA captures intra- and inter-packet dependencies, efficiently handling the challenges of modeling token-to-token relations across packets. Table 1 shows that ET-BERT, using only token-level attention, performs worse. As shown in Table 4, results on CrossPlatform(Android) shows removing flow attention raises error rates by 16.4%, and using token-mean embeddings for packets lowers accuracy to 62.19%, underscoring the importance of both attentions.

Impact of the Number of Packets. Figure 5 shows the impact of packet count. On the left, we observe that as more packets provide additional information, the F1 score

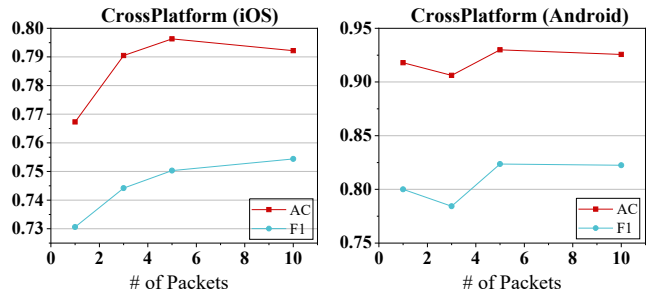


Figure 5: Impact of the Number of Packets.

Methods	CrossPlatform (Android)		CrossPlatform (iOS)	
	AC	F1	AC	F1
header only	78.77%	64.98%	52.96%	41.66%
payload only	72.46%	65.80%	63.82%	55.30%
All	93.00%	82.36%	79.63%	75.03%

Table 5: Impact of the Component of Packets.

increases, confirming our expectations. However, on the CrossPlatform (Android) dataset, using just one packet outperforms using three, with 91.79% accuracy, higher than all baseline models using five packets. This suggests that in some datasets, the initial packets may contain the most critical information, and adding more packets may not always improve performance. If packet relationships are not effectively modeled, it could even harm performance. In this case, the reason our model did not perform well may be due to the discrepancy between the pre-training stage, where five packets were used, and the fine-tuning stage, where only three were used, leading to a mismatch in distribution.

Impact of the Component of Packets. Table 5 highlights the advantages of payload-based methods in traffic classification. When using only the packet header, the model's performance is significantly lower, showing that the header lacks sufficient information for accurate classification. However, when focusing on the payload, which contains the actual data, the model's F1-scores improve, especially on complex datasets like Cross Platform (iOS). This demonstrates that payload-based methods are more effective in capturing the essential characteristics of traffic, making them superior for encrypted traffic classification. Combining both header and payload yields the best results.

Conclusion

We introduced MIETT to address challenges in encrypted traffic classification. Through novel model architecture and pre-training strategies, MIETT effectively learns temporal and flow-specific dynamics. Experiments show that MIETT outperforms existing methods across five datasets. Future work will explore traffic classification in resource-constrained scenarios (Zhou 2024).

Acknowledgments

This work is partially supported by the National Science and Technology Major Project (2022ZD0114805), NSFC (62376118), Key Program of Jiangsu Science Foundation (BK20243012), Collaborative Innovation Center of Novel Software Technology and Industrialization.

References

- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Draper-Gil, G.; Lashkari, A. H.; Mamun, M. S. I.; and Ghorbani, A. A. 2016. Characterization of encrypted and vpn traffic using time-related. In *Proceedings of the 2nd international conference on information systems security and privacy (ICISSP)*, 407–414.
- He, H. Y.; Yang, Z. G.; and Chen, X. N. 2020. PERT: Payload Encoding Representation from Transformer for Encrypted Traffic Classification. In *2020 ITU Kaleidoscope: Industry-Driven Digital Transformation, Kaleidoscope, Ha Noi, Vietnam, December 7-11, 2020*, 1–8. IEEE.
- Korczyński, M.; and Duda, A. 2012. Classifying service flows in the encrypted skype traffic. In *2012 IEEE International Conference on Communications (ICC)*, 1064–1068. IEEE.
- Lashkari, A. H.; Gil, G. D.; Mamun, M. S. I.; and Ghorbani, A. A. 2017. Characterization of tor traffic using time based features. In *International Conference on Information Systems Security and Privacy*, volume 2, 253–262. SciTePress.
- Lin, K.; Xu, X.; and Gao, H. 2021. TSCRNN: A novel classification scheme of encrypted traffic based on flow spatiotemporal features for efficient management of IIoT. *Computer Networks*, 190: 107974.
- Lin, X.; Xiong, G.; Gou, G.; Li, Z.; Shi, J.; and Yu, J. 2022. ET-BERT: A Contextualized Datagram Representation with Pre-training Transformers for Encrypted Traffic Classification. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, 633–642. ACM.
- Liu, C.; He, L.; Xiong, G.; Cao, Z.; and Li, Z. 2019. Fs-net: A flow sequence network for encrypted traffic classification. In *IEEE INFOCOM 2019-IEEE Conference On Computer Communications*, 1171–1179. IEEE.
- Lotfollahi, M.; Jafari Siavoshani, M.; Shirali Hossein Zade, R.; and Saberian, M. 2020. Deep packet: A novel approach for encrypted traffic classification using deep learning. *Soft Computing*, 24(3): 1999–2012.
- Moore, A. W.; and Papagiannaki, K. 2005. Toward the accurate identification of network applications. In *International workshop on passive and active network measurement*, 41–54. Springer.
- Neto, E. C. P.; Dadkhah, S.; Ferreira, R.; Zohourian, A.; Lu, R.; and Ghorbani, A. A. 2023. CIIoT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment. *Sensors*, 23(13): 5941.
- Qian, C.; Li, X.; Wang, Q.; Zhou, G.; and Shao, H. 2024. NetBench: A Large-Scale and Comprehensive Network Traffic Benchmark Dataset for Foundation Models. In *2024 IEEE International Workshop on Foundation Models for Cyber-Physical Systems & Internet of Things (FMSys)*, 20–25. IEEE Computer Society.
- Van Ede, T.; Bortolameotti, R.; Continella, A.; Ren, J.; Dubois, D. J.; Lindorfer, M.; Choffnes, D.; Van Steen, M.; and Peter, A. 2020. Flowprint: Semi-supervised mobile-app fingerprinting on encrypted network traffic. In *Network and distributed system security symposium (NDSS)*, volume 27.
- Wang, P.; Ye, F.; Chen, X.; and Qian, Y. 2018. Datanet: Deep learning based encrypted network traffic classification in sdn home gateway. *IEEE Access*, 6: 55380–55391.
- Wang, Y.; Zhang, Z.; Guo, L.; and Li, S. 2011. Using entropy to classify traffic more deeply. In *2011 IEEE Sixth International Conference on Networking, Architecture, and Storage*, 45–52. IEEE.
- Yao, H.; Liu, C.; Zhang, P.; Wu, S.; Jiang, C.; and Yu, S. 2019. Identification of encrypted traffic through attention mechanism based long short term memory. *IEEE transactions on big data*, 8(1): 241–252.
- Zhao, R.; Zhan, M.; Deng, X.; Wang, Y.; Wang, Y.; Gui, G.; and Xue, Z. 2023. Yet Another Traffic Classifier: A Masked Autoencoder Based Traffic Transformer with Multi-Level Flow Representation. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023*, 5420–5427. AAAI Press.
- Zhou, Z.-H. 2024. Learnability with time-sharing computational resource concerns. *National Science Review*, 11(10): nwae204.