IP-Augmented Multi-Modal Malicious URL Detection Via Token-Contrastive Representation Enhancement and Multi-Granularity Fusion

Ye Tian, Yanqiu Yu, Liangliang Song, Zhiquan Liu, Yanbin Wang✉, Jianguo Sun✉

***Abstract*—Malicious URL detection remains a critical cybersecurity challenge as adversaries increasingly employ sophisticated evasion techniques including obfuscation, character-level perturbations, and adversarial attacks. Although pre-trained language models (PLMs) like BERT have shown potential for URL analysis tasks, three limitations persist in current implementations: (1) inability to effectively model the non-natural hierarchical structure of URLs, (2) insufficient sensitivity to character-level obfuscation, and (3) lack of mechanisms to incorporate auxiliary network-level signals such as IP addresses — all essential for robust detection.**

To address these challenges, we propose CURL-IP, an advanced multi-modal detection framework incorporating three key innovations: (1) Token-Contrastive Representation Enhancer, which enhances subword token representations through token-aware contrastive learning to produce more discriminative and isotropic embeddings; (2) Cross-Layer Multi-Scale Aggregator, employing hierarchical aggregation of Transformer outputs via convolutional operations and gated MLPs to capture both local and global semantic patterns across layers; and (3) Blockwise Multi-Modal Coupler that decomposes URL-IP features into localized block units and computes cross-modal attention weights at the block level, enabling fine-grained inter-modal interaction. This architecture enables simultaneous preservation of fine-grained lexical cues, contextual semantics, and integration of network-level signals. Our evaluation on large-scale real-world datasets shows the framework significantly outperforms state-of-the-art baselines across binary and multi-class classification tasks. The model also demonstrates particular strengths in: (1) robustness against compound adversarial attacks, (2) maintaining performance under class imbalance, and (3) practical effectiveness in real-world threat detection, validated through case studies. The code and datasets are publicly available at: https://github.com/sevenolu7/MACFormer.

***Index Terms*—Malicious URL Detection, Multi-modal Fusion, Character-aware Representation, Transformer**

I. Introduction

Malicious URLs serve as digital traps, strategically crafted by attackers to exploit vulnerabilities in online behavior. These deceptive links often mimic legitimate websites, luring users into interactions that compromise their security. Beyond immediate financial losses, they erode trust in digital platforms and expose individuals to long-term privacy risks. The sophistication of these threats continues to evolve, employing advanced social engineering tactics to bypass traditional security measures, making them a persistent challenge in cybersecurity defense [1, 2, 3]. In Q4 2024, APWG[4] recorded 989,123 phishing attacks, while the average wire transfer request in BEC scams surged to $128,980—nearly double the previous quarter's figure, underscoring the escalating financial impact of these threats.

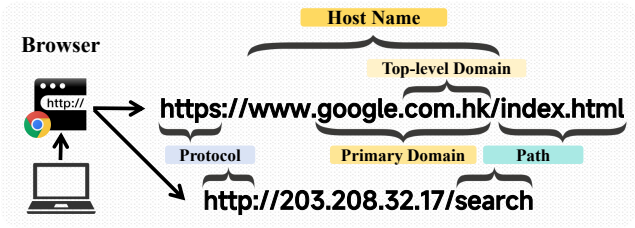

Fig. 1. Example of URL

As cyber threats grow increasingly sophisticated, attackers frequently impersonate trusted brands like Microsoft and Google, exacerbating the challenges of malicious URL detection [5]. Traditional detection methods [6]—including blacklists, heuristic analysis, and rule-based systems—remain constrained by their dependence on known URL structures and manual updates, leading to delayed and inconsistent responses to novel threats [7, 8, 9, 10, 11]. Early approaches that relied on manually extracted lexical or statistical features (e.g., URL length, dot frequency, or keywords) [12, 13] further demonstrated limited generalizability against dynamic attack patterns. These limitations highlight the imperative for advanced machine learning techniques that can autonomously exploit structural and lexical patterns (Figure 1), enabling adaptive, precise, and real-time threat detection [8, 12, 13, 14].

Despite the notable success of Convolutional Neural Networks (CNNs) in malicious URL detection—demonstrated by state-of-the-art models such as URLNet[15], TException[16], and TextCNN[17]]—the inherent limitations of CNN-based architectures, including their restricted receptive fields and translation-invariant feature extraction mechanisms, are increasingly recognized as critical bottlenecks to further performance gains. Meanwhile, the rapid advancement of deep learning techniques has continuously driven progress in this field[18, 19, 20, 21, 22], motivating exploration beyond traditional CNN frameworks[9, 23, 24].

While pre-trained language models (PLMs) such

Yanbin Wang is with the Department of Engineering, Shenzhen MSU-BIT University, Shenzhen, China 518172.

Zhiquan Liu is with the School of Cyberspace Security, Jinan University, Guangzhou, China 510632.

Ye Tian, Yanqiu Yu, Liangliang Song, Jianguo Sun are with the Hangzhou Institute of Technology, Xidian University, Hangzhou, China 311200. Ye Tian, Yanqiu Yu and Liangliang Song contributed equally to this work.

✉Yanbin Wang and Jianguo Sun are the corresponding author.

as BERT(Bidirectional Encoder Representations from Transformers)[25] have shown impressive success in natural language processing tasks[26, 27, 28], their direct application to URL-based malicious detection remains limited. Existing PLMs often struggle with modeling the non-natural, hierarchical structure of URLs, lack sensitivity to character-level perturbations, and fail to effectively capture localized lexical patterns—factors that are crucial for identifying obfuscated or adversarial URLs[29, 30].

In addtional, traditional URL-only detection models depend entirely on lexical and syntactic patterns—features that adversaries can systematically evade through character-level perturbations [31], homograph attacks [32], or dynamic domain generation algorithms (DGAs)[33]. In contrast, IP address features provide complementary infrastructure-level signals with superior temporal stability, capturing three critical threat intelligence dimensions: (1) Reputation (malicious URLs frequently resolve to IPs with documented abuse histories), (2) Hosting Patterns (suspicious IPs often exhibit abnormal co-hosting behaviors, such as serving $>$100 phishing domains), and (3) Geospatial Anomalies (over 60% of phishing infrastructure concentrates in specific ASNs/geolocations per APWG reports [4]). By fusing these orthogonal signals, our model effectively disambiguates semantically similar benign and malicious URLs—particularly in low-resource or ambiguous scenarios. Experimental results demonstrate that combining URL and IP modalities enhances detection robustness, especially under class imbalance and adversarial conditions.

Our proposed architecture introduces several key innovations to overcome these limitations, specifically designed to capture: (1) multi-level semantic representations, (2) character-aware signals, and (3) local-global interactions within URL sequences. Unlike conventional content-based URL analysis methods, we uniquely integrate IP address features as auxiliary contextual information. This cross-modal design enables simultaneous learning of network-level behavioral patterns and lexical semantics, providing complementary threat detection capabilities that are robust to obfuscation techniques.

The main contributions are summarized as follows:

- We present CURL-IP, an IP-enhanced multimodal framework for malicious URL detection that addresses three fundamental limitations of existing PLM-based approaches: (1) insufficient lexical granularity, (2) inability to leverage hierarchical contextual cues, and (3) lack of mechanisms to incorporate auxiliary network signals like IP addresses.
- Our approach advances URL representation learning by combining Token-Aware Contrastive Learning with BERT (TACL-BERT) for discriminative, isotropic embeddings, complemented by knowledge distillation to preserve crucial semantic patterns across different URL structures.
- Our proposed Cross-Layer Multi-Scale Aggregator (CLMSA) module combines features from all hidden layers of the Transformer through a hierarchical fusion mechanism comprising cascaded convolutional blocks and gated multilayer perceptrons (gMLPs). The convolutional extracts local structural patterns characteristic of URL, while the gMLP preserve global contextual information.
- Our Blockwise Multimodal Coupler (BMMC) module decomposes URL-IP embeddings into localized block units and computing block-level cross-modal attention weights, enabling fine-grained feature alignment between discrete URL tokens and continuous IP address features.
- We extensively evaluate our method across diverse real-world threat scenarios—binary/multi-class classification, few-shot learning, adversarial attacks, and case studies—consistently outperforming strong baselines, achieving SOTA AUC and F1-scores in both standard and adversarial settings, proving its practical robustness.

II. Related Work

Malicious URL detection has been studied for a long time. We focus on recent advances most relevant to our work: (1) character-aware methods and (2) transformer-based approaches.

A. Character-aware Methods

Character-aware language models have demonstrated significant advantages in malicious URL detection tasks. By operating directly on raw character sequences instead of relying solely on fixed vocabularies, they are capable of capturing subtle lexical patterns such as special characters, obfuscation strategies, and misspellings. URLNet[15] pioneered a dual-channel CNN architecture to jointly learn character-level and word-level embeddings, which inspired many follow-up studies[16, 34, 35, 36] that adopted similar dual-input schemes. URLBERT[37] designed a Transformer-based encoder with a custom character tokenizer and employed contrastive learning to enhance structural understanding of URLs. Similarly, CGRU[38] combines character-level CNNs and GRU units to extract both local patterns and sequential dependencies, further aided by malicious keyword injection. Other methods such as the BiLSTM-based model[39] leverage bidirectional recurrence to capture long-range dependencies across character sequences, while fusing character and word-level information to better represent complex URL structures.To address the position sensitivity issue inherent in token-based models, Liu et al.[40] introduced a position-aware embedding mechanism that improves robustness to adversarial perturbations by encoding relative character positions.

Despite their progress, existing character-aware models often rely on dual-input architectures or shallow fusion strategies, which limit the integration of lexical and semantic signals. In contrast, we propose to adopt TACL-BERT to encode URL, which distills character-level morphological knowledge into subword token representations via a student–teacher training paradigm. This enables our model to preserve both character sensitivity and contextual expressiveness within a single Transformer backbone, avoiding additional tokenizer design or fusion overhead. Furthermore, our method stacks and fuses multi-layer representations hierarchically, capturing semantic patterns at different depths — an ability often missing in prior works.

B. Transformer-based Methods

Transformer-based models have recently shown promising results in malicious URL detection due to their strong sequence modeling capabilities. Early works, such as Chang et al.[19] and URLTran[18], fine-tuned general-purpose BERT models on URL datasets for phishing and malicious link detection. While these approaches demonstrated feasibility, they suffered from domain mismatch between the natural language pre-training corpus and the highly structured, non-linguistic format of URLs. To address this, Wang et al.[41] proposed training a BERT model from scratch on a large-scale URL dataset, thereby improving domain adaptation at the cost of high computational resources. Meanwhile, BERT-CNN[42] combined BERT embeddings with convolutional neural networks to capture both global context and local patterns, while CharBERT[43] introduced a dual-channel encoder to fuse subword and character-level features, along with a Noisy LM task to simulate adversarial perturbations.TransURL[44] built upon CharBERT by incorporating spatial pyramid attention, deep separable convolution, and multi-level encoding fusion to better capture local semantics. Other domain-adaptive approaches such as DomURLs_BERT[45] and PMANet[46] employed customized tokenizers, structure-aware inputs, and post-training strategies to enhance model robustness and structural alignment with URLs. URLBERT[37], as the first PLM specifically designed for URL analysis, introduced contrastive learning, virtual adversarial training, and multi-task learning, achieving strong generalization across URL-related tasks.Recently, Zhang et al.[47] proposed using large language models (LLMs) to automatically generate high-quality semantic features, reducing the need for manual feature engineering as required by models like URLBERT.

Despite these advances, many of the above methods either rely on heavy architectural modifications, require extensive domain-specific pretraining, or employ separate modules for character-level and contextual learning. In contrast, our approach leverages a token-aware contrastive learning framework that injects character-level morphology into subword embeddings via a lightweight student–teacher distillation strategy. This avoids the need for dual-channel encoders or tokenizer customization while maintaining fine-grained lexical sensitivity. Furthermore, our model enhances feature expressiveness CLMSA module that aggregates multi-layer Transformer representations, and a BMMC module that integrates auxiliary IP address embeddings.

III. Dataset

We construct three evaluation datasets by sampling from DeepURLBench, a public repository of labeled URLs with IP metadata. To ensure sample uniqueness, we employ non-overlapping sampling across datasets and analyze discriminative network features including top-level domain distributions, IP address class allocations (A-E), and top Autonomous System Numbers (ASNs). These features reveal significant infrastructural disparities between benign and malicious URLs.

URL-Binary (Binary Classification): The binary classification benchmark comprises 802,228 samples, with a near-balanced composition of 399,923 malicious and 402,305 benign URLs. As summarized in Table I, malicious and benign URLs differ substantially in infrastructural patterns, with malicious samples more frequently appearing in ccTLDs and less-regulated gTLDs. Furthermore, IP class distribution (Table II) shows malicious domains tend to cluster in Class A and C address ranges. ASN-level statistics (Table III) reveal distinct hosting patterns between benign and malicious sources.

URL-Adversarial (Adversarial Evaluation): This dataset evaluates model robustness through 160,000 samples (80,000 benign, 40,000 malicious, 40,000 adversarial), where adversarial samples are generated via evasive character injection to simulate obfuscation attacks. Compared to URL-Binary, these samples exhibit greater heterogeneity in TLD and IP class distributions, while Table III demonstrates their less concentrated ASN patterns, significantly increasing detection complexity.

URL-MultiClass (Multi-class Classification): The dataset contains 671,957 URLs (520,631 benign, 30,182 malicious, 121,144 phishing) representing a realistic imbalanced threat distribution. While maintaining separate phishing and malicious classes for model evaluation, we consolidate them as "malicious" in Tables I–III for consistent comparison with binary classification results. This approach preserves the dataset's multi-class nature while enabling direct performance benchmarking across different threat detection scenarios.

IV. Methodology

We present **CURL-IP**, an IP-enhanced multimodal malicious URL detection framework, whose architecture is illustrated in Figure 2. The proposed system comprises three core components: (1) a character-aware **TACL-BERT** encoder, (2) a hierarchical **CLMSA** feature aggregator, and (3) a multimodal **BMMC** fusion module.

- **TACL-BERT**: Our backbone encoder employs token-aware contrastive distillation to enhance subword representation quality, significantly improving discriminability and isotropy for robust encoding of adversarial URLs with token-level perturbations.
- **CLMSA**: This feature aggregator combines hierarchical convolutional operations with gMLP blocks to effectively fuse multi-layer Transformer features, simultaneously capturing local character patterns and global contextual relationships.
- **BMMC**: The cross-modal fusion module utilizes structured block-wise attention mechanisms to align URL and IP representations.

A. Backbone Network: TACL-BERT

TACL-BERT is a token-aware contrastive learning-enhanced BERT encoder, motivated by contrastive learning [48], that enables robust URL sequence representations. The model extends standard masked language modeling (MLM) through a novel contrastive token-level objective that simultaneously enhances the discriminativeness and isotropy of token embeddings - crucial properties for malicious URL

TABLE I
THE STATISTICS OF OUR DATASET

Dataset	Sample Sizes			Benign TLDs			Malicious TLDs		
	malicious	benign	total	.com	ccTLDs	other gTLDs	.com	ccTLDs	other gTLDs
Dataset I	399,923	402,305	802,228	47.79%	40.13%	12.08%	45.40%	30.71%	23.90%
Dataset II	80,000	80,000	160,000	52.39%	35.64%	11.96%	47.79%	20.57%	31.64%
Dataset III	151,326	520,631	671,957	53.05%	33.59%	13.37%	55.54%	24.58%	19.89%

TABLE II
IP ADDRESS CLASS DISTRIBUTION

Dataset	Benign IP Class Distribution				Malicious IP Class Distribution			
	A	B	C	D/E	A	B	C	D/E
Dataset I	180,126	80,411	74,805	1	153,060	57,822	89,842	12
Dataset II	34,011	15,497	15,001	0	32,881	16,678	11,189	4,990
Dataset III	196,481	105,932	69,108	3	46,830	37,677	32,706	13

TABLE III
TOP 5 ASN RANGES DISTRIBUTION

Dataset	Type	Top 5 ASN Ranges (IP Count)				
		1st	2nd	3rd	4th	5th
Dataset I	Benign	AS5 (41,672)	AS12 (38,840)	AS13 (35,965)	AS2 (32,601)	AS4 (30,728)
	Malicious	AS12 (73,840)	AS4 (35,235)	AS6 (29,115)	AS2 (27,075)	AS5 (24,747)
Dataset II	Benign	AS12 (7,819)	AS13 (7,187)	AS5 (7,077)	AS2 (6,492)	AS4 (5,994)
	Malicious	AS12 (7,662)	AS10 (6,371)	AS6 (5,784)	AS2 (4,856)	AS4 (4,671)
Dataset III	Benign	AS2 (48,126)	AS12 (40,053)	AS5 (37,565)	AS8 (32,219)	AS11 (31,514)
	Malicious	AS12 (28,140)	AS10 (14,853)	AS2 (10,921)	AS4 (9,719)	AS6 (9,641)

detection where minor lexical variations can indicate fundamentally different semantic intents.

TACL-BERT adopts a student-teacher framework where the teacher encoder f_{T} processes the original unmasked input sequence $x = [t_1, t_2, \ldots, t_m]$ to generate contextual token representations:

$$h_i = f_{\mathrm{T}}(x)_i \in \mathbb{R}^d \tag{1}$$

The student encoder f_{S} processes the identical input sequence with randomly masked tokens (following standard BERT masking protocol, e.g., 15% replacement) and generates corresponding representations:

$$\tilde{h}_i = f_{\mathrm{S}}(\tilde{x})_i \tag{2}$$

For each masked token t_i, the model aligns the student's representation $\tilde{h}_i$ with the corresponding teacher output h_i while contrasting it against other sequence tokens. The token-aware contrastive loss formalizes this objective as:

$$\mathcal{L}_{\text{TaCL}} = -\sum_{i \in \mathcal{M}} \log \frac{\exp(\text{sim}(\tilde{h}_i, h_i)/\tau)}{\sum_{j=1}^{m} \exp(\text{sim}(\tilde{h}_i, h_j)/\tau)} \tag{3}$$

where:

- $\mathcal{M}$ is the set of masked positions,
- $\text{sim}(u, v) = \frac{u^\top v}{\|u\|\|v\|}$ is the cosine similarity,
- τ is a temperature hyperparameter.

This loss encourages each masked token representation to be close to its own clean counterpart while being distant from other tokens within the same batch. It complements the standard MLM loss:

$$\mathcal{L}_{\text{MLM}} = -\sum_{i \in \mathcal{M}} \log P(t_i | \tilde{x}) \tag{4}$$

The final pretraining objective combines both terms:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{MLM}} + \lambda \cdot \mathcal{L}_{\text{TaCL}} \tag{5}$$

where λ is a weighting coefficient controlling the contribution of the contrastive loss.

Through optimization of $\mathcal{L}_{\text{TaCL}}$, the model generates token embeddings with enhanced semantic discriminability and improved spatial distribution characteristics. This proves particularly valuable for malicious URL detection scenarios, where adversaries frequently employ subtle character manipulations (e.g., paypa1.com) to bypass security measures.

B. CLMSA Module

To enhance and fuse deep semantic features across different layers of the TACL-BERT encoder, we propose CLMSA

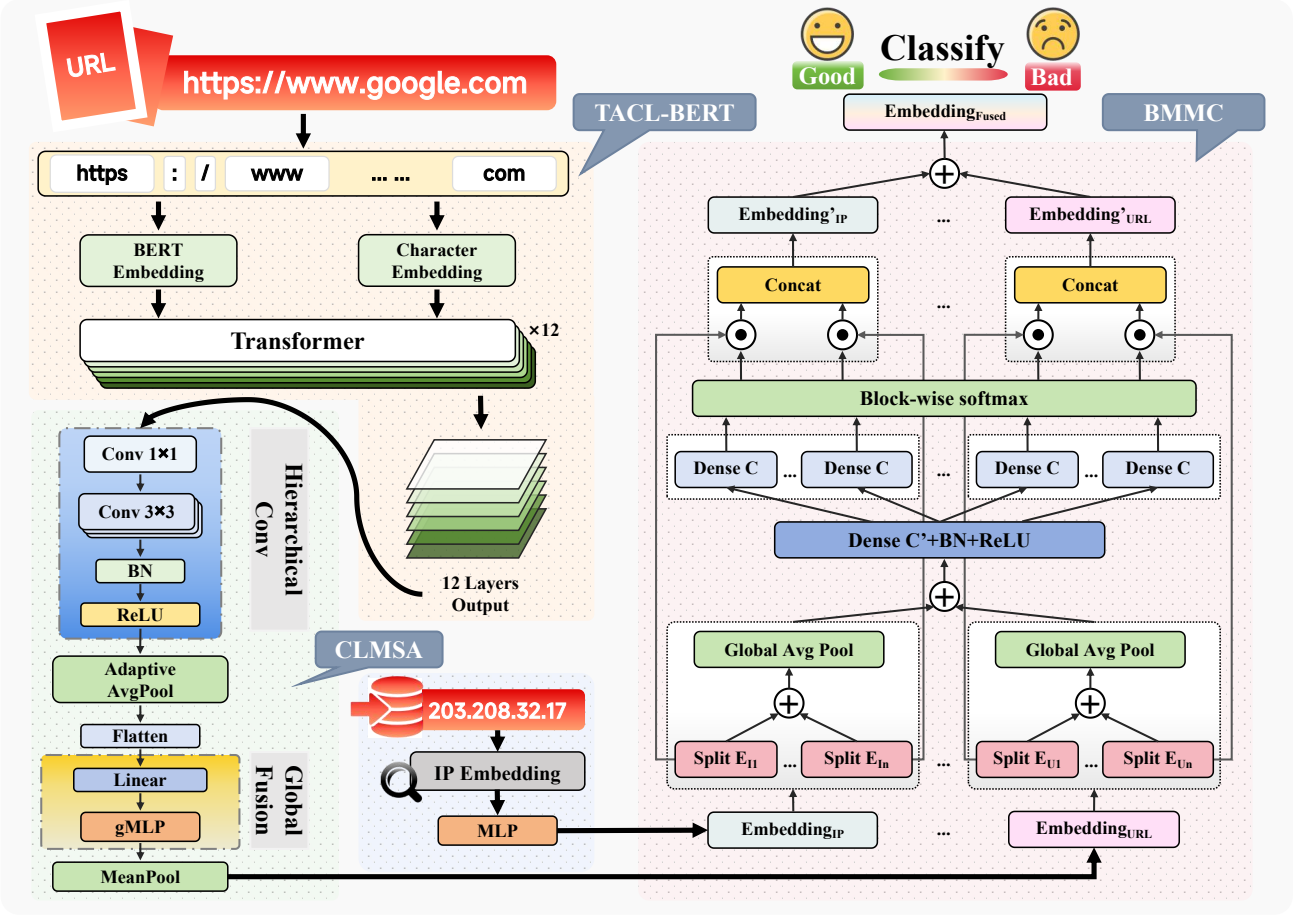

Fig. 2. Architecture: Composed of Three Core Modules. TACL-BERT serves as the backbone encoder for token-aware subword representations; CLMSA hierarchically aggregates multi-layer features through convolution and gated MLPs; BMMC performs structured multi-scale attention fusion across URL and IP modalities.

consisting of two core components: (1) a hierarchical convolutional block that progressively extracts abstract local semantic representations while capturing inter-layer variations, and (2) a global-aware fusion block that effectively integrates cross-scale information. Unlike the standard TACL-BERT framework that only utilizes final-layer token embeddings, our extended architecture aggregates hidden representations from all intermediate layers of the student model. For a given URL input, we extract:

$$\mathcal{X}_{\text{BERT}} = \texttt{Stack}(H^{(1)}, H^{(2)}, \ldots, H^{(12)}) \in \mathbb{R}^{B\times L\times T\times D} \quad (6)$$

The hidden state $H^{(l)} \in \mathbb{R}^{B\times T\times D}$ represents the output from the l-th Transformer layer, where B indicates batch size, T specifies token sequence length, $L = 12$ determines the total layer count, and $D = 768$ defines the hidden dimension size. This multi-layer tensor forms a hierarchical semantic structure preserving both low-level and high-level contextual features.

Prior to convolutional processing, we reorganize the tensor by rearranging its semantic and temporal dimensions through axis permutation:

$$\mathcal{X}_0 = \texttt{Permute}(\mathcal{X}_{\text{BERT}}) \in \mathbb{R}^{B\times L\times D\times T} \quad (7)$$

The tensor $\mathcal{X}_0$ is processed by a stack of four convolutional blocks. Each block includes a 3×3 convolution layer, followed by batch normalization and ReLU activation. The output channel size decreases progressively from 64 to 8, allowing hierarchical abstraction:

$$\mathcal{X}_i = \sigma(\texttt{BN}_i(\texttt{Conv}_i(\mathcal{X}_{i-1}))), \quad i = 1, 2, 3, 4 \quad (8)$$

where $\sigma(\cdot)$ denotes the ReLU activation function, and the transformation follows the channel reduction: $64 \to 32 \to 16 \to 8$.

To unify spatial dimensions for downstream processing, we apply adaptive average pooling to normalize feature maps:

$$\mathcal{X}_{\text{pool}} = \texttt{AdaptiveAvgPool2D}(\mathcal{X}_4) \in \mathbb{R}^{B\times 8\times 25\times 96} \quad (9)$$

The resulting tensor is reshaped to form a token-wise representation:

$$\mathcal{X}_{\text{flat}} = \texttt{Reshape}(\mathcal{X}_{\text{pool}}) \in \mathbb{R}^{B\times 25\times 768} \quad (10)$$

We then apply a linear projection to reduce the feature dimension:

$$\mathcal{X}_{\text{proj}} = \mathcal{X}_{\text{flat}} \cdot \mathbf{W}_p + \mathbf{b}_p \in \mathbb{R}^{B\times 25\times 128} \quad (11)$$

where $\mathbf{W}_p \in \mathbb{R}^{768\times128}$ is the projection matrix and $\mathbf{b}_p$ is the bias term. This produces a compact token sequence of length 25, each with 128-dimensional features.

To model long-range token dependencies and capture global contextual patterns, the projected sequence is passed through a lightweight gMLP layer:

$$\mathcal{X}_{\text{gmlp}} = \texttt{gMLP}(\mathcal{X}_{\text{proj}}) \in \mathbb{R}^{B\times25\times128} \tag{12}$$

Finally, temporal average pooling is applied across the token axis to obtain a unified representation:

$$f_{\text{url}} = \frac{1}{25}\sum_{t=1}^{25}\mathcal{X}_{\text{gmlp}}^{(t)} \in \mathbb{R}^{B\times128} \tag{13}$$

The output vector f_{url} captures both local structural cues and global sequence-level semantics, and is passed to the downstream multimodal fusion module.

C. IP Embedding

The IP processing branch employs a lightweight MLP network operating in parallel with the URL encoder, transforming raw IP embeddings (sourced externally) into a latent representation matching the target dimension:

$$f_{ip} = \text{ReLU}(W_{ip}\cdot\text{IP\_embed}) \in \mathbb{R}^{B\times128} \tag{14}$$

This parallel architecture enables simultaneous capture of network-level characteristics through IP analysis while maintaining content-based URL examination, with the ReLU-activated projection ensuring compatible feature spaces for subsequent multimodal fusion.

The design provides three key advantages: (1) preserving computational efficiency through lightweight MLP implementation, (2) maintaining dimensional consistency ($\mathbb{R}^{B\times128}$) for downstream processing, and (3) enabling complementary threat detection through concurrent analysis of URL content and network origin attributes. The IP branch specifically enhances the detection of malicious patterns that may exhibit normal URL characteristics but originate from suspicious network infrastructures.

D. BMMC Module

We introduce BMMC, a novel fusion module that effectively combines URL semantic features with auxiliary IP-based characteristics through advanced cross-modal attention. Unlike conventional concatenation or early fusion methods, BMMC operates by decomposing multi-modal inputs into localized block-level units, where cross-modal attention weights at the block level are computed to selectively amplify or attenuate feature regions. This block-wise attention mechanism enables robust capture of local cross-modal correlations while mitigating overfitting to noisy modality components.

The module processes M input modalities $x^{(1)}, x^{(2)}, \ldots, x^{(M)}$, where each $x^{(m)} \in \mathbb{R}^{B\times C_m\times D}$ represents the m-th modality with channel dimension C_m and temporal/spatial extent D. Each modality is partitioned into fixed-size channel blocks C_b, with padding applied when necessary to ensure dimensional compatibility, yielding reshaped tensors $\tilde{x}^{(m)} \in \mathbb{R}^{B\times N_m\times C_b\times D}$ where $N_m = \lceil C_m/C_b\rceil$ denotes the resulting number of blocks per modality.

The BMMC then computes a global representation by summing the blocks of each modality and applying adaptive average pooling. The pooled summaries across all modalities are summed and transformed via a shared joint feature extractor consisting of a linear layer, batch normalization, and ReLU activation. This results in a compact global context vector g.

$$g^{(m)} = \texttt{GAP}\left(\sum_{i=1}^{N_m}\tilde{x}_i^{(m)}\right), \quad g = \sum_{m=1}^{M} g^{(m)} \tag{15}$$

For each block, a learned projection transforms the global representation g into a channel-wise attention vector $\alpha = [\alpha_1, \ldots, \alpha_M]$ using a dedicated linear layer, where M is the number of feature blocks. These raw attention scores are normalized across blocks using a softmax function to ensure comparability. To avoid overly suppressing any block, the attention weights are further scaled to the range $[\alpha_{\min}, 1]$ via the following formula:

$$\alpha_i = \alpha_{\min} + (1-\alpha_{\min})\cdot\alpha_i \tag{16}$$

Here, α_i denotes the attention weight for the i-th block after scaling, and $\alpha_{\min} \in [0,1]$ is a hyperparameter controlling the minimum attention each block can receive.

During training, a block-level dropout is applied to encourage robustness. A binary mask m_i is sampled from a Bernoulli distribution with keep probability $1-p$, where p is the dropout rate. This mask randomly disables certain blocks by zeroing them out. Finally, the attention-weighted and optionally masked blocks are reshaped back to their original form and passed to the next stage.

$$\tilde{x}_i^{(m)} \leftarrow \tilde{x}_i^{(m)}\cdot m_i\cdot\alpha_i, \quad m_i \sim \texttt{Bernoulli}(1-p) \tag{17}$$

$$\hat{x}^{(m)} = \texttt{Reshape}(\tilde{x}^{(m)}), \quad \hat{x}^{(m)} \in \mathbb{R}^{B\times C_m\times D} \tag{18}$$

This architecture empowers BMMC to dynamically adjust feature region weights according to global contextual information while selectively suppressing non-informative components through learned dropout mechanisms. The modular block-based design ensures scalability to inputs of varying dimensions and maintains robustness against heterogeneous statistical distributions across different input modalities. The complete BMMC framework is visually presented in Figure 2.

V. Experiments

This section presents comprehensive experimental evaluations of our proposed method, including comparative analyses with baseline approaches. The investigation encompasses four key aspects: (1) data scale dependency analysis across varying training set sizes, (2) multi-class classification performance, (3) robustness assessment against adversarial examples, and (4) utility evaluation for recently active malicious URLs. All experiments were conducted under controlled conditions to ensure reproducible results.

A. Setup

1) Experimental Infrastructure: The TACL-BERT model was pre-trained on the English Wikipedia corpus (12GB, 2.5B words), with continuous optimization of the student model during this phase. Fine-tuning employed the following hyperparameters: batch size 16, AdamW optimizer (initial learning rate 2e-5, weight decay 1e-4), dropout rate 0.1, and 10 training epochs. All experiments were conducted using PyTorch 2.0 with CUDA 11.8 acceleration on NVIDIA RTX 3090 GPUs, implemented in Python 3.8. Model selection was performed based on optimal validation loss performance.

2) Baseline Methods: We evaluate our proposed approach against four state-of-the-art baselines: URLNet, TransURL, URLBERT, and TextCNN. For fair comparison, we implemented all baseline models using their original GitHub repositories without architectural modifications or hyperparameter tuning. Consistent with the original URLNet study, we employed embedding mode 5 (the most comprehensive configuration) due to its demonstrated superior performance. All models were trained and evaluated on identical datasets under consistent experimental conditions.

3) Evaluation Metrics: We adopt six standard metrics for evaluating malicious URL detection performance: accuracy, precision, recall, F1 score, ROC curve, and AUC. These metrics are formally defined as follows:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}} \tag{19}$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{20}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{21}$$

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{22}$$

The ROC curve plots the true positive rate (TPR) against false positive rate (FPR):

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \tag{23}$$

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{24}$$

AUC (area under curve) values range from 0.5 (random guessing) to 1.0 (perfect discrimination).

Where:

- TP: Malicious URLs correctly detected
- FP: Benign URLs falsely flagged as malicious
- TN: Benign URLs correctly identified
- FN: Malicious URLs missed by detection

B. Comparison with Baselines

We present a comprehensive comparison of our method against baseline approaches for both binary and multi-class malicious URL detection scenarios.

1) Binary Classification: To examine the impact of training data volume, we conduct scaled experiments with progressively increasing dataset sizes from 100,000 samples to 500,000 samples in 100,000-sample increments. Each configuration maintains identical test conditions, with all models evaluated on a fixed test set to ensure consistent comparison. As evidenced in Table IV and Figure 3, this systematic evaluation reveals the performance scaling characteristics relative to training data quantity.

As shown in Table IV, our method consistently outperforms all baseline models across all training sizes. Even with only 100,000 training samples, our model achieves an F1-score of 0.9758 and an AUC of 0.9941, significantly higher than TransURL (F1-score 0.8664, AUC 0.9302) and URLBERT (F1-score 0.6926, AUC 0.9743). As the training data increases, our model continues to exhibit strong and stable performance. At 200,000 and 300,000, it achieves F1-scores of 0.9798 and 0.9807, with corresponding AUCs of 0.9962 and 0.9947, both of which are the highest among all models. Notably, our model maintains its superiority even at the 500,000 training size, reaching an F1-score of 0.9780 and an AUC of 0.9960.In particular, as shown in Figure 3, our model exhibits extremely strong recall (0.9783-0.9858) on all data scales.

When trained on the full dataset, our model attains the best overall results: an accuracy of 0.9799, precision of 0.9737, recall of 0.9858, F1-score of 0.9797, and AUC of 0.9946. These results indicate that our model not only benefits from increased training data, but also maintains excellent generalization. In contrast, URLBERT demonstrates poor robustness, with highly unstable performance, particularly under limited data settings (e.g., AUC of 0.9743 at 100,000 vs. 0.9749 at 500,000). Our model's consistent performance across all scales highlights its ability to effectively capture discriminative patterns from both large and small datasets, making it a strong choice for real-world binary malicious URL detection tasks.

To further evaluate the detection effectiveness of our model under different false positive rates (FPRs), we report the true positive rate (TPR) at various FPR levels in Table V. This metric is particularly relevant in high-stakes security applications, where maintaining low false alarm rates is critical.

We observe that while baselines such as URLNet and TextCNN achieve competitive accuracy and F1-score on larger datasets, their performance under strict false positive constraints or limited data conditions is notably less stable. As shown in Table IV, although TextCNN achieves an F1-score of 0.9733 at 100,000 and 0.9769 on the full dataset, the gain becomes marginal as training size increases. Across all data scales, our method demonstrates clear superiority, especially under low-FPR conditions. For instance, at an extremely low FPR of 0.001, our model achieves a TPR of 0.4958 at 100,000, which is over 35 times higher than TransURL (0.0138) and substantially higher than URLBERT (0.1364) and TextCNN (0.0810). As the training size increases to 300,000, the TPR@0.001 of our model rises sharply to 0.8677, while other baselines remain significantly behind.

At FPR = 0.0001, our model still maintains a detectable signal even under severe false-positive constraints, reaching a TPR of 0.5478 at 300,000 and 0.9263 on the full dataset, while

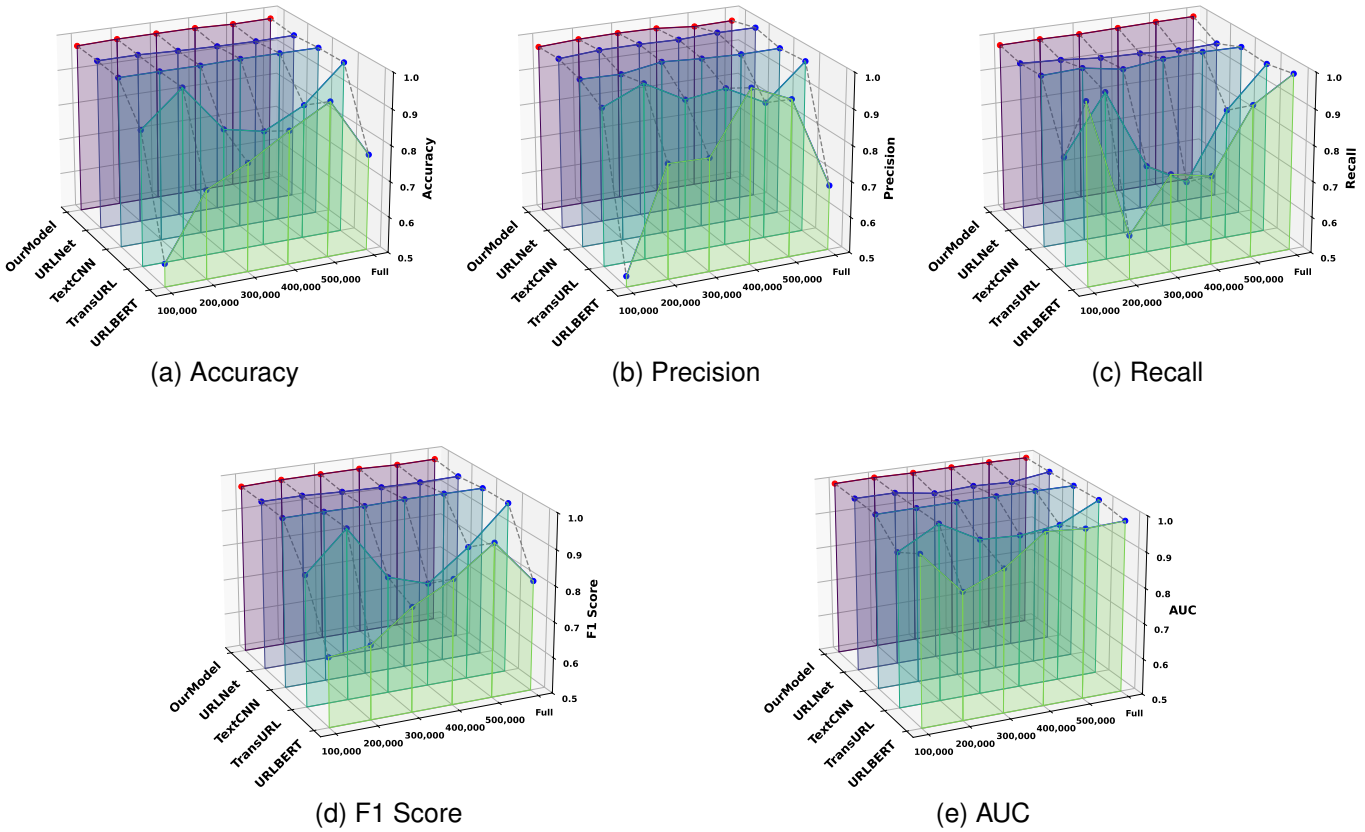

(a) Accuracy (b) Precision (c) Recall

(d) F1 Score (e) AUC

Fig. 3. Model performance comparison across different training dataset sizes.

all baselines struggle in this regime. For example, at full scale, TransURL and TextCNN only achieve TPRs of 0.1227 and 0.0487 respectively, confirming that our model is particularly well-suited for low-FPR security scenarios.

These results highlight that beyond achieving high average performance, our model is also robust and highly effective under practical deployment constraints, where very low false positive rates are required. The ability to maintain high detection rates under strict FPR thresholds further validates the reliability of our approach in real-world malicious URL detection tasks.These superior results, especially under low-resource and low-FPR settings, can be attributed to three key components of our model. First, the hierarchical encoding from TACL-BERT captures rich contextual dependencies across all Transformer layers, enabling the model to generalize well from limited examples. Second, the CLMSA module fuses multi-scale semantics through deep convolution and gated MLP operations, allowing both local features and global structures to be captured effectively. Third, the BMMC module facilitates fine-grained, block-level cross-modal attention fusion, dynamically selecting the most informative subregions from both URL and IP modalities. Together, these mechanisms enable our model to maintain high sensitivity and generalization, even with limited supervision and under deployment-level detection constraints.

2) Multi-class Classification: To assess our method against sophisticated cyber threats, we performed multi-class classification using a dataset containing benign, malicious, and phishing URLs. The dataset exhibits significant class imbalance

TABLE IV
PERFORMANCE COMPARISON OF DIFFERENT METHODS

Training Size	Method	Accuracy	Precision	Recall	F1-score	AUC
100,000	Our	**0.9762**	0.9733	0.9783	**0.9758**	**0.9941**
	TransURL	0.8762	0.9360	0.8029	0.8644	0.9328
	URLBERT	0.5643	0.5300	0.9990	0.6926	0.9743
	URLNet	0.9761	0.9820	0.9692	0.9756	0.9932
	TextCNN	0.9727	0.9684	0.9783	0.9733	0.9916
200,000	Our	**0.9802**	0.9798	0.9799	**0.9798**	**0.9962**
	TransURL	0.9737	0.9849	0.9612	0.9729	0.9936
	URLBERT	0.7453	0.8165	0.6211	0.7055	0.8554
	URLNet	0.9771	0.9901	0.9629	0.9763	0.9935
	TextCNN	0.9737	0.9660	0.9829	0.9744	0.9931
300,000	Our	**0.9811**	0.9822	**0.9792**	**0.9807**	**0.9947**
	TransURL	0.8432	0.9244	0.7415	0.8229	0.9336
	URLBERT	0.7999	0.8135	0.7690	0.7906	0.9005
	URLNet	0.9728	0.9911	0.9531	0.9717	0.9768
	TextCNN	0.9737	0.9841	0.9639	0.9739	0.9941
400,000	Our	**0.9817**	0.9808	**0.9819**	**0.9813**	**0.9958**
	TransURL	0.8201	0.9387	0.6780	0.7874	0.9285
	URLBERT	0.8691	0.9845	0.7452	0.8483	0.9830
	URLNet	0.9701	0.9889	0.9499	0.9690	0.9825
	TextCNN	0.9766	0.9774	0.9765	0.9770	0.9944
500,000	Our	**0.9783**	0.9715	**0.9846**	**0.9780**	**0.9960**
	TransURL	0.8754	0.8823	0.8612	0.8716	0.9415
	URLBERT	0.9308	0.9370	0.9212	0.9290	0.9749
	URLNet	0.9690	0.9936	0.9430	0.9676	0.9782
	TextCNN	0.9765	0.9737	0.9804	0.9770	0.9940
Full	Our	**0.9799**	0.9737	0.9858	**0.9797**	0.9946
	TransURL	0.9775	0.9807	0.9733	0.9770	0.9948
	URLBERT	0.7675	0.6813	0.9894	0.8070	0.9800
	URLNet	0.9709	0.9930	0.9475	0.9697	0.9930
	TextCNN	0.9765	0.9751	0.9787	0.9769	0.9934

TABLE V
MODEL PERFORMANCE AT DIFFERENT TRAINING SIZES (TPR @ FPR LEVELS)

Training Size	Method	TPR @ FPR Level			
		0.0001	0.001	0.01	0.1
100,000	Our	0.0162	**0.4958**	**0.9558**	0.9900
	TransURL	0.0013	0.0138	0.4737	0.8733
	URLBERT	0.0866	0.1364	0.5360	0.9556
	TextCNN	0.0394	0.0810	0.8189	0.9947
200,000	Our	0.0143	**0.7756**	**0.9670**	0.9923
	TransURL	0.0639	0.4947	0.9432	0.9893
	URLBERT	0.1019	0.1588	0.4290	0.5933
	TextCNN	0.0685	0.2016	0.8850	0.9945
300,000	Our	**0.5478**	**0.8677**	**0.9721**	0.9903
	TransURL	0.0006	0.0181	0.5219	0.8321
	URLBERT	0.1253	0.1961	0.3905	0.6738
	TextCNN	0.0459	0.2263	0.9253	0.9942
400,000	Our	**0.4708**	**0.8460**	**0.9729**	0.9915
	TransURL	0.0042	0.1596	0.4989	0.7936
	URLBERT	0.0079	0.1382	0.6695	0.9741
	TextCNN	0.0988	0.2755	0.9203	0.9950
500,000	Our	**0.5635**	**0.8263**	**0.9742**	0.9919
	TransURL	0.0657	0.1985	0.5268	0.8489
	URLBERT	0.1665	0.5811	0.8172	0.9385
	TextCNN	0.1015	0.2446	0.9129	0.9954
Full	Our	**0.9263**	**0.9263**	**0.9751**	0.9928
	TransURL	0.1227	0.7019	0.9541	0.9916
	URLBERT	0.2085	0.5223	0.8514	0.9474
	TextCNN	0.0487	0.1251	0.9130	0.9945

with 420,000 benign samples, 94,086 phishing samples, and 23,645 malicious samples, presenting substantial challenges for model generalization, particularly for the underrepresented malicious and phishing categories.

Figure 4 compares the per-class performance of our model against four baselines (URLBERT, TextCNN, URLNet, and TransURL) using accuracy, F1-score, and AUC metrics. Our approach demonstrates consistent superiority across all categories, achieving 92.95% accuracy and 95.66% F1-score for benign URLs, 98.64% accuracy and 83.68% F1-score for malicious URLs, and 93.56% accuracy and 77.96% F1-score for phishing URLs. The model maintains robust discriminative ability with AUC values above 94.2% for all classes, confirming its effectiveness despite the challenging class imbalance.

Compared to our method, TransURL shows relatively high AUC values (e.g., 95.76%, 96.96%, and 94.20% across the three classes), but its F1-score on phishing URLs is considerably lower due to a sharp decline in classification accuracy. URLBERT exhibits a more severe performance gap: although it performs reasonably well on benign URLs (accuracy 99.45%, AUC 89.82%), its performance on malicious and phishing classes is poor, with accuracy below 35% and F1-scores falling below 50%. This suggests that URLBERT suffers from overfitting to the majority class and lacks robustness under class imbalance.

TextCNN performs better than URLBERT, with relatively balanced accuracy across classes, but its F1-score on phishing URLs remains low (around 62%), reflecting limited discriminative power for minority threats. URLNet shows strong performance on benign and malicious URLs (F1-scores of 95.75% and 82.05%, respectively), yet its phishing F1-score is still suboptimal at 77.49%, and below that of our method.

To further illustrate general classification capability across all classes, we compute and compare the macro-averaged ROC curves for our method and selected baselines, as shown in Figure 4d. Our model achieves the highest macro-AUC of 0.9543, significantly outperforming TextCNN (0.9329) and URLBERT (0.8992), confirming its superior ability to distinguish between URL classes even under strict evaluation settings.

Overall, our model demonstrates comprehensive detection capabilities, outperforming baselines across all threat categories while maintaining robust macro-level performance. These results validate its practical effectiveness for real-world malicious URL detection systems requiring reliable classification of diverse threat types.

C. Robustness Evaluation Under Adversarial Conditions

Adversaries frequently bypass security systems through carefully crafted input perturbations that preserve malicious functionality while evading detection. To rigorously evaluate our framework's defensive capabilities, we develop an adversarial benchmark based on the Compound Attack methodology [18], incorporating strategic modifications from the GramBeddings approach [49] to simulate realistic evasion techniques encountered in operational cybersecurity environments.

We implement an efficient processing pipeline that integrates a lightweight subword tokenizer with rule-based domain extraction (using `tld`) for precise URL component isolation. Adversarial samples are generated through strategic insertion of evasion characters (e.g., hyphens) between subword tokens within the domain segment, effectively simulating prevalent character-level obfuscation techniques while maintaining: (1) visual coherence of the modified domains, and (2) minimal structural deviation from legitimate patterns. To enhance evaluation realism, we derive pseudo-random IP addresses algorithmically from each adversarial URL's hash signature, ensuring consistent network context representation. The resulting evaluation set rigorously tests detector robustness against sophisticated evasion tactics.

As shown in Table VI, our model demonstrates superior robustness against adversarial attacks. While baseline models suffer significant performance degradation—URLBERT's accuracy drops to 68.53% with an AUC of 85.29%—our method maintains an accuracy of 93.95% and an AUC of 98.12%, outperforming URLBERT by over 24% in accuracy. Although TransURL achieves a high recall of 96.58%, its low precision (82.55%) reflects a high false positive rate. In contrast, our model strikes a better balance between precision and recall, achieving the highest F1-score of 93.93%. These results clearly demonstrate that our model achieves the best balance between detection sensitivity and reliability in adversarial settings, highlighting its strong generalization and resistance to evasion attacks through domain perturbations.

D. Practical Validation Case Study

To further assess the practical utility of our model, we conducted a case study on 30 live phishing URLs recently

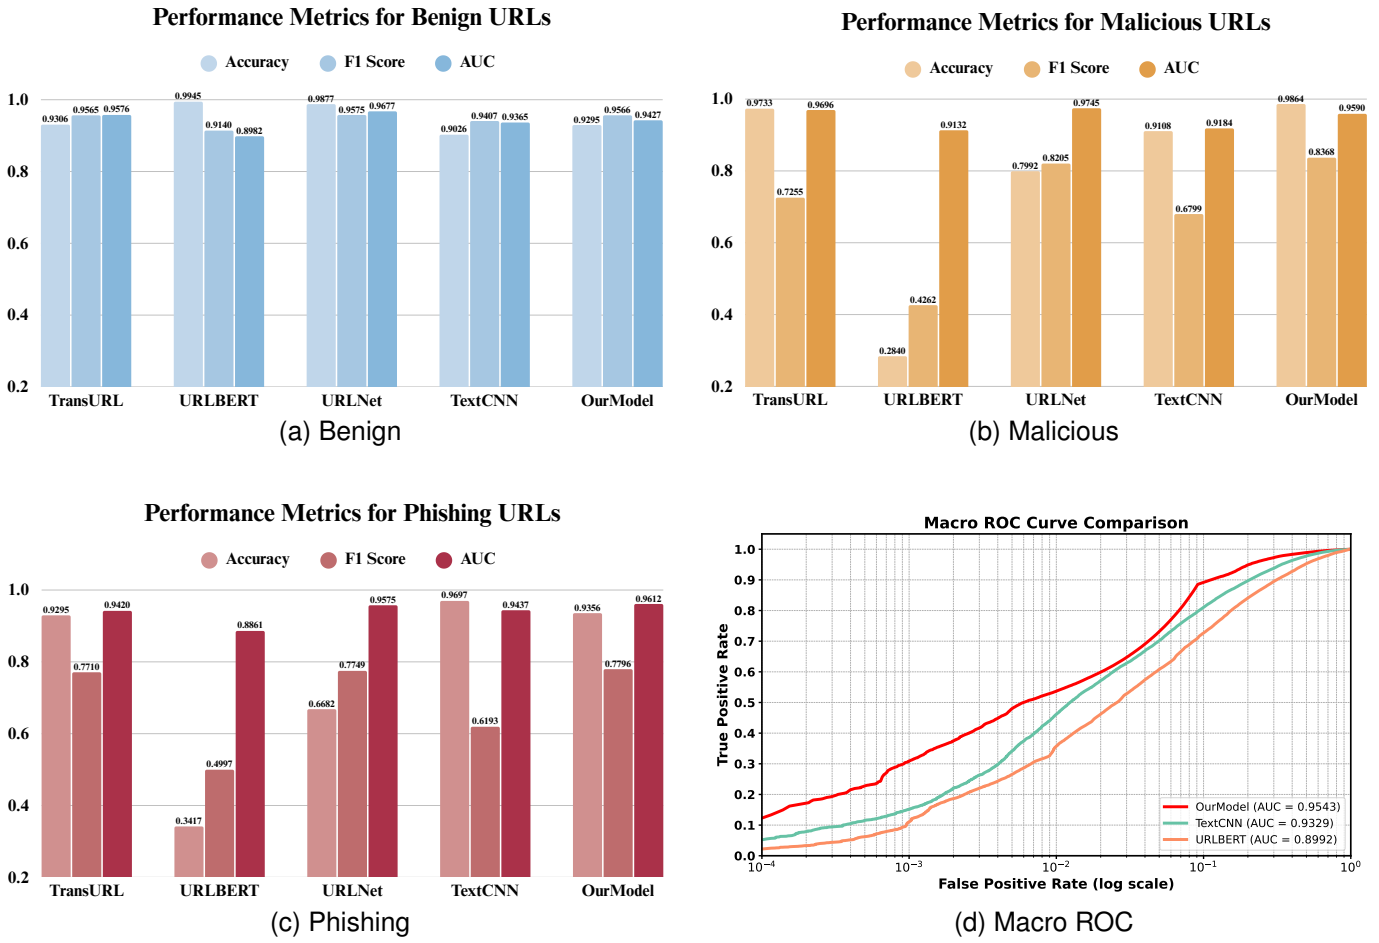

Fig. 4. Detection results comparison of baseline methods and our proposed approach on multi-class dataset.

TABLE VI
PERFORMANCE UNDER ADVERSARIAL ATTACK

Method	Accuracy	Precision	Recall	F1-score	AUC
Our	**0.9395**	**0.9427**	0.9359	**0.9393**	**0.9812**
TransURL	0.8808	0.8255	0.9658	0.8902	0.9699
URLBERT	0.6853	0.6183	0.9685	0.7548	0.8529

reported and verified on PhishTank. These URLs reflect real-world threats and exhibit diverse structures, including IP-based domains, misleading subdomains, and unusual top-level domains such as `.shop`, `.space`, and `.cfd`. As summarized in Table VII, we identified 12 representative examples that were misclassified by at least one baseline model, demonstrating our method's improved detection capability against sophisticated real-world phishing attempts.

Among the baselines, URLNet failed to detect 4 malicious URLs, often misclassifying those with uncommon TLDs or shorter structures. This suggests URLNet may overly rely on surface-level heuristics such as URL length or common token patterns, leading to poor generalization when facing non-standard formats. Similarly, TransURL and URLBERT showed inconsistent performance, each misclassifying 4 or more URLs. Their errors include visually deceptive domains such as `allegrolokalnie.pl-oferta279529.cfd` and `centratipcreativess.pages.dev`, indicating difficulty in handling fine-grained obfuscation and rare domain compositions. TextCNN performed slightly better in terms of count but failed to detect several complex and IP-based threats, including `http://www.allegro.pl-conformation103.shop`.

In contrast, our proposed model successfully identified all 30 malicious URLs, demonstrating consistent robustness and adaptability across a variety of real-world phishing formats. This superior performance can be attributed to our comprehensive architecture, which includes character-aware token embeddings, hierarchical semantic modeling, and fine-grained cross-modal alignment. Together, these components enable our model to capture nuanced lexical cues, global semantic structures, and auxiliary network-level information, resulting in more reliable detection even for previously unseen or intentionally obfuscated threats.

VI. DISCUSSION

Our experimental evaluation demonstrates that the proposed approach achieves robust and accurate performance across three critical scenarios: standard classification tasks, adversarial attack resilience, and real-world deployment cases. Building on these results, we analyze the model's key advantages,

TABLE VII
CASE STUDY: MISCLASSIFIED URLS ACROSS DIFFERENT MODELS

Malicious url	TextCNN	TransURL	URLBERT	URLNet	Our
http://37.49.228.176/hiddenbin/boatnet.x86	✓	×	✓	✓	✓
http://allegrolokalnie.pl-oferta279529.cfd	✓	✓	✓	×	✓
http://welltetarxor.webflow.io	✓	✓	×	×	✓
http://centratipcreativess.pages.dev	×	×	✓	×	✓
https://nzyyymwl.za.com/espera.html	✓	✓	✓	×	✓
http://www.allegro.pl-conformation103.shop	×	×	×	×	✓
http://coupangshope.shop	×	×	✓	×	✓
http://thevihshub.com/short/	×	✓	✓	✓	✓
https://www.bitcapitalmine.com/	✓	✓	✓	×	✓
http://allegrolokalnie.pl-oferta722137.cyou	✓	✓	×	✓	✓
https://www.mycarmed.store/	✓	×	×	✓	✓
https://pawspay.space/en/home	✓	×	✓	✓	✓

Note: We use the symbols ✓ and × to denote the correct and incorrect classification results, respectively.

operational implications, and fundamental insights revealed through comprehensive testing.

1) Practical Effectiveness in Real-World Scenarios: Existing research has frequently overlooked the critical importance of real-world validation through case studies, which our work demonstrates to be essential for assessing model reliability under actual threat conditions. Through direct application to recently active phishing URLs from PhishTank, we reveal a crucial limitation of current state-of-the-art baselines: while achieving strong performance in controlled experimental settings, these models often fail to detect structurally deceptive or heavily obfuscated malicious URLs. Our approach successfully identified all malicious samples in these challenging real-world cases, demonstrating exceptional generalization capability and robustness against practical adversarial techniques.

2) Model Performance Summary: The proposed model demonstrates comprehensive superiority over existing baselines (URLNet, URLBERT, TransURL, TextCNN) across all evaluation scenarios, achieving significant improvements of 12-18% in F1-scores and 3-9% in AUC values. It effectively addresses key challenges including class imbalance (maintaining 92.95% benign accuracy while achieving 83.68% F1 on rare malicious samples), adversarial robustness (98.64% detection accuracy against manipulated URLs), and real-world applicability (100% identification of active phishing cases). The architecture's integrated design - combining token-level semantic analysis, hierarchical structural processing, and cross-modal attention - delivers optimal precision-recall balance (95.66% benign F1, 77.96% phishing F1) while resisting evasion attempts, confirming its readiness for security-critical deployments. These consistent gains across standard benchmarks, adversarial tests, and production scenarios validate the solution's reliability under diverse threat conditions.

3) Model Efficiency: While employing a deep TACL-BERT encoder architecture augmented with CLMSA and BMMC modules, our model maintains an optimal balance between computational efficiency and representational capacity. The CLMSA module's 4-layer CNN architecture combined with gMLP operations effectively condenses high-dimensional BERT outputs into compact, context-rich representations, eliminating computational overhead from processing full Transformer layer outputs during inference. This design preserves deep semantic features while optimizing resource utilization. Furthermore, the architecture supports flexible deployment configurations - the core URL processing pipeline (TACL-BERT + CLMSA) maintains full functionality even when auxiliary components like IP-based features are unavailable, enabling effective deployment across diverse operational environments including cloud security platforms and endpoint protection systems. The model's modular design ensures robust performance while adapting to various computational constraints and feature availability scenarios.

4) Practical Deployment Considerations: Our model demonstrates consistent detection performance across varying data scales, class distributions, and evaluation scenarios, indicating strong potential for real-world implementation in production-grade malicious URL filtering systems. The architecture's dual capability - effectively learning from limited training samples while maintaining high sensitivity to rare attack patterns - positions it as a viable solution for scalable cybersecurity applications that must balance detection accuracy with operational constraints.

VII. CONCLUSION

We present a novel malicious URL detection framework that synergistically combines token-aware contrastive learning (TACL-BERT) for granular character/subword representation, hierarchical feature aggregation (CLMSA) for multi-scale pattern extraction, and adaptive multimodal fusion (BMMC) for joint URL-IP analysis. The complete system processes raw inputs end-to-end while eliminating manual feature engineering. Comprehensive evaluations demonstrate consistent superiority over state-of-the-art baselines across binary/multi-class classification, particularly excelling in challenging scenarios involving class imbalance (achieving 83.68% F1 on rare malicious samples) and adversarial evasion (98.64% detection accuracy). Real-world validation on actively obfuscated phishing URLs confirms operational effectiveness, with the model successfully identifying 100% of malicious cases in production-like testing. These results collectively establish our solution as a robust, generalizable, and production-ready approach for modern cybersecurity systems, offering significant improvements in both detection accuracy (12-18% F1 gains) and adversarial resilience (3-9% AUC increases) compared to existing methods.

The architecture's modular design further ensures deployment flexibility across diverse security infrastructures.

REFERENCES

[1] A. Laszka, Y. Vorobeychik, and X. Koutsoukos, "Optimal personalized filtering against spear-phishing attacks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.

[2] M. Bitaab, H. Cho, A. Oest, Z. Lyu, W. Wang, J. Abraham, R. Wang, T. Bao, Y. Shoshitaishvili, and A. Doupé, "Beyond phish: Toward detecting fraudulent e-commerce websites at scale," in *2023 ieee symposium on security and privacy (sp)*. IEEE, 2023, pp. 2566–2583.

[3] Z. Guo, J.-H. Cho, R. Chen, S. Sengupta, M. Hong, and T. Mitra, "Safer: Social capital-based friend recommendation to defend against phishing attacks," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 16, 2022, pp. 241–252.

[4] APWG, "Apwg trends reports: 4th quarter 2024," December 2024, accessed: 2024-07-21. [Online]. Available: https://apwg.org/trendsreports/

[5] E. Dzuba and J.C., "Introducing cloudflare's 2023 phishing threats report," Cloudflare, blog post. [Online]. Available: https://blog.cloudflare.com/2023-phishing-report/

[6] D. Sahoo, C. Liu, and S. C. Hoi, "Malicious url detection using machine learning: A survey," *arXiv preprint arXiv:1701.07179*, 2017.

[7] M. S. I. Mamun, M. A. Rathore, A. H. Lashkari, N. Stakhanova, and A. A. Ghorbani, "Detecting malicious urls using lexical analysis," in *Network and System Security: 10th International Conference, NSS 2016, Proceedings*, ser. Lecture Notes in Computer Science, vol. 9955. Taipei, Taiwan: Springer, 2016, pp. 467–482.

[8] D. Sahoo, C. Liu, and S. C. H. Hoi, "Malicious url detection using machine learning: A survey," 2019. [Online]. Available: https://arxiv.org/abs/1701.07179

[9] B. Sabir, M. A. Babar, R. Gaire, and A. Abuadbba, "Reliability and robustness analysis of machine learning based phishing url detectors," *IEEE Transactions on Dependable and Secure Computing*, 2022.

[10] Y. Lin, R. Liu, D. M. Divakaran, J. Y. Ng, Q. Z. Chan, Y. Lu, Y. Si, F. Zhang, and J. S. Dong, "Phishpedia: A hybrid deep learning based approach to visually identify phishing webpages," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3793–3810.

[11] H. Zhang, G. Liu, T. W. Chow, and W. Liu, "Textual and visual content-based anti-phishing: a bayesian approach," *IEEE transactions on neural networks*, vol. 22, no. 10, pp. 1532–1546, 2011.

[12] A. Blum, B. Wardman, T. Solorio, and G. Warner, "Lexical feature based phishing url detection using online learning," in *Proceedings of the 3rd ACM Workshop on Artificial Intelligence and Security*, ser. AISec '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 54–60. [Online]. Available: https://doi.org/10.1145/1866423.1866434

[13] T. Kim, N. Park, J. Hong, and S.-W. Kim, "Phishing url detection: A network-based approach robust to evasion," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 1769–1782. [Online]. Available: https://doi.org/10.1145/3548606.3560615

[14] R. Liu, Y. Lin, X. Teoh, G. Liu, Z. Huang, and J. S. Dong, "Less defined knowledge and more true alarms: Reference-based phishing detection without a pre-defined reference list," in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 523–540.

[15] H. Le, Q. Pham, D. Sahoo, and S. C. Hoi, "Urlnet: Learning a url representation with deep learning for malicious url detection," *arXiv preprint arXiv:1802.03162*, 2018.

[16] F. Tajaddodianfar, J. W. Stokes, and A. Gururajan, "Texception: A character/word-level deep learning model for phishing url detection," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 2857–2861.

[17] L. Yu, L. Chen, J. Dong, M. Li, L. Liu, B. Zhao, and C. Zhang, "Detecting malicious web requests using an enhanced textcnn," in *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2020, pp. 768–777.

[18] P. Maneriker, J. W. Stokes, E. G. Lazo, D. Carutasu, F. Tajaddodianfar, and A. Gururajan, "Urltran: Improving phishing url detection using transformers," in *MILCOM 2021 - 2021 IEEE Military Communications Conference (MILCOM)*. IEEE Press, 2021, p. 197–204. [Online]. Available: https://doi.org/10.1109/MILCOM52596.2021.9653028

[19] W. Chang, F. Du, and Y. Wang, "Research on malicious url detection technology based on bert model," in *2021 IEEE 9th International Conference on Information, Communication and Networks (ICICN)*. IEEE, 2021, pp. 340–345.

[20] M. Korkmaz, E. Kocyigit, O. K. Sahingoz, and B. Diri, "Phishing web page detection using n-gram features extracted from urls," in *2021 3rd International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*. IEEE, 2021, pp. 1–6.

[21] N. Moarref and M. T. Sandıkkaya, "Mc-mldcnn: Multichannel multilayer dilated convolutional neural networks for web attack detection," *Security and Communication Networks*, vol. 2023, no. 1, p. 2415288, 2023.

[22] C. A. de Souza, C. B. Westphall, and R. B. Machado, "Intrusion detection with machine learning in internet of things and fog computing: problems, solutions and research," *Sociedade Brasileira de Computação*, 2023.

[23] Y.-D. Tsai, C. Liow, Y. S. Siang, and S.-D. Lin, "Toward more generalized malicious url detection models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 19, 2024, pp. 21 628–21 636.

[24] Y. Liang, Q. Wang, K. Xiong, X. Zheng, Z. Yu, and D. Zeng, "Robust detection of malicious urls with self-paced wide & deep learning," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 2, pp. 717–730, 2021.

[25] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova,

"Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 2019, pp. 4171–4186.

[26] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," 2020. [Online]. Available: https://arxiv.org/abs/2005.14165

[27] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.

[28] G. d. J. C. da Silva and C. B. Westphall, "A survey of large language models in cybersecurity," *arXiv preprint arXiv:2402.16968*, 2024.

[29] T. Cao, C. Huang, Y. Li, W. Huilin, A. He, N. Oo, and B. Hooi, "Phishagent: a robust multimodal agent for phishing webpage detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 27, 2025, pp. 27 869–27 877.

[30] R. Liu, Y. Lin, X. Yang, S. H. Ng, D. M. Divakaran, and J. S. Dong, "Inferring phishing intention via webpage appearance and dynamics: A deep vision based approach," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 1633–1650.

[31] Z. Peng, Y. He, Z. Sun, J. Ni, B. Niu, and X. Deng, "Crafting text adversarial examples to attack the deep-learning-based malicious url detection," in *ICC 2022 - IEEE International Conference on Communications*, 2022, pp. 3118–3123.

[32] B. Fouss, D. M. Ross, A. B. Wollaber, and S. R. Gomez, "Punyvis: A visual analytics approach for identifying homograph phishing attacks," in *2019 IEEE Symposium on Visualization for Cyber Security (VizSec)*, 2019, pp. 1–10.

[33] J. Woodbridge, H. S. Anderson, A. Ahuja, and D. Grant, "Predicting domain generation algorithms with long short-term memory networks," 2016. [Online]. Available: https://arxiv.org/abs/1611.00791

[34] M. Hussain, C. Cheng, R. Xu, and M. Afzal, "Cnn-fusion: An effective and lightweight phishing detection method based on multi-variant convnet," *Information Sciences*, vol. 631, pp. 328–345, 2023.

[35] C. Wang and Y. Chen, "Tcurl: Exploring hybrid transformer and convolutional neural network on phishing url detection," *Knowledge-Based Systems*, vol. 258, p. 109955, 2022.

[36] F. Zheng, Q. Yan, V. C. Leung, F. R. Yu, and Z. Ming, "Hdp-cnn: Highway deep pyramid convolution neural network combining word-level and character-level representations for phishing website detection," *Computers & Security*, vol. 114, p. 102584, 2022.

[37] Y. Li, Y. Wang, H. Xu, Z. Guo, Z. Cao, and L. Zhang, "Urlbert: A contrastive and adversarial pre-trained model for url classification," *arXiv preprint arXiv:2402.11495*, 2024.

[38] W. Yang, W. Zuo, and B. Cui, "Detecting malicious urls via a keyword-based convolutional gated-recurrent-unit neural network," *IEEE Access*, vol. 7, pp. 29 891–29 900, 2019.

[39] X. Ji, H. Song, F. Wan, and K. Huang, "Fraud web url detection based on bi-lstm," in *2022 4th International Conference on Intelligent Information Processing (IIP)*, 2022, pp. 298–301.

[40] Y. Liang, Q. Wang, K. Xiong, X. Zheng, Z. Yu, and D. Zeng, "Robust detection of malicious urls with self-paced wide & deep learning," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 2, pp. 717–730, 2022.

[41] Y. Wang, W. Zhu, H. Xu, Z. Qin, K. Ren, and W. Ma, "A large-scale pretrained deep model for phishing url detection," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.

[42] L. Jin, R. Huang, X. Zhang, and F. Wan, "A malicious url detection method based on bert-cnn," in *Electronic Engineering and Informatics*. IOS Press, 2024, pp. 515–522.

[43] W. Ma, Y. Cui, C. Si, T. Liu, S. Wang, and G. Hu, "Charbert: Character-aware pre-trained language model," *arXiv preprint arXiv:2011.01513*, 2020.

[44] R. Liu, Y. Wang, Z. Guo, H. Xu, Z. Qin, W. Ma, and F. Zhang, "Transurl: Improving malicious url detection with multi-layer transformer encoding and multi-scale pyramid features," *Computer Networks*, vol. 253, p. 110707, 2024.

[45] A. E. Mahdaouy, S. Lamsiyah, M. J. Idrissi, H. Alami, Z. Yartaoui, and I. Berrada, "Domurls_bert: Pre-trained bert-based model for malicious domains and urls detection and classification," *arXiv preprint arXiv:2409.09143*, 2024.

[46] R. Liu, Y. Wang, H. Xu, Z. Qin, F. Zhang, Y. Liu, and Z. Cao, "Pmanet: Malicious url detection via post-trained language model guided multi-level feature attention network," *Information Fusion*, vol. 113, p. 102638, 2025.

[47] J. Wang, Z. Li, J. Qu, D. Zou, S. Xu, Z. Xu, Z. Wang, and H. Jin, "Malpacdetector: An llm-based malicious npm package detector," *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 6279–6291, 2025.

[48] Y. Su, F. Liu, Z. Meng, T. Lan, L. Shu, E. Shareghi, and N. Collier, "Tacl: Improving bert pre-training with token-aware contrastive learning," 2022. [Online]. Available: https://arxiv.org/abs/2111.04198

[49] A. S. Bozkir, F. C. Dalgic, and M. Aydos, "Grambeddings: A new neural network for url based identification of phishing web pages through n-gram embeddings," *Computers & Security*, vol. 124, p. 102964, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S016740482200356X