ISSN: 2395-7852

International Journal of Advanced Research in Arts, Science, Engineering & Management (IJARASEM)

Volume 6, Issue 6, November 2019

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

IMPACT FACTOR: 5.454

| www.ijarasem.com | ijarasem@gmail.com | +91-9940572462 |

International Journal of Advanced Research in Arts, Science, Engineering & Management (IJARASEM)

| ISSN: 2395-7852 | www.ijarasem.com | | Impact Factor: 5.454 |Bimonthly, Peer Reviewed & Referred Journal|

| Volume 6, Issue 6, November 2019 |

Deep Packet Inspection using Convolutional Neural Networks for Encrypted Traffic Classification

Mohamed Elhoseny

Faculty of Computers and Information, Mansoura University, Egypt

ABSTRACT: With the proliferation of HTTPS and virtual private networks (VPNs), traditional network traffic monitoring techniques that rely on payload inspection are rapidly losing efficacy. In enterprise and security contexts, identifying applications generating encrypted traffic—without violating privacy or decrypting content—has become a critical need. This paper explores the use of convolutional neural networks (CNNs) to classify encrypted network traffic by analyzing packet-level flow features such as size, inter-arrival timing, and directionality, while completely bypassing payload inspection. We collect and label encrypted traffic from popular applications such as Skype, YouTube, Dropbox, and Facebook, and convert these sessions into flow-based representations suitable for CNN input. Our model achieves over 89% classification accuracy, significantly outperforming baseline methods like k-Nearest Neighbors (k-NN) and Support Vector Machines (SVM). We also develop a lightweight, real-time monitoring module using TensorFlow Lite and validate its integration within a passive network sensor. The proposed approach maintains user privacy, enables compliance auditing, and facilitates dynamic policy enforcement. While the method is challenged by obfuscated traffic and application mimicry, it represents a meaningful step toward privacy-preserving encrypted traffic analytics. Future work includes anomaly detection, online learning, and extension to zero-day application behaviors.

KEYWORDS: encrypted traffic classification, CNN, deep packet inspection, flow analysis, network monitoring, privacy-preserving analytics, TensorFlow, real-time detection

I. INTRODUCTION

In the past, network security relied heavily on Deep Packet Inspection (DPI), which enabled firewalls, intrusion detection systems (IDS), and monitoring tools to examine packet payloads for signatures or behavioral patterns. However, the near-universal adoption of TLS (Transport Layer Security), HTTPS, and VPN tunneling has rendered traditional DPI increasingly ineffective. With over 90% of web traffic now encrypted, network administrators face a dilemma: how to monitor and classify traffic without decrypting it and infringing on user privacy?

This challenge is particularly acute in enterprise networks, where administrators must enforce compliance with application usage policies—such as preventing unauthorized file sharing or bandwidth-heavy streaming—without exposing sensitive personal data. Encrypted traffic classification becomes a privacy-preserving alternative to DPI, where traffic is inferred based on flow-level statistics rather than decrypted content.

Recent advances in machine learning, particularly deep learning, have enabled new approaches to this problem. Convolutional Neural Networks (CNNs), initially popularized in image classification, have shown promise in identifying spatial and temporal patterns in network flows. By encoding traffic sequences as matrices or tensors, CNNs can extract hierarchical features that are robust to noise, encryption, and obfuscation.

This paper investigates the application of CNNs to the problem of encrypted traffic classification. We build and train a deep learning model using labeled traffic from applications such as Skype, YouTube, Dropbox, and Facebook. Using only packet sizes, timing intervals, and direction metadata, we achieve high classification accuracy without violating content confidentiality. The model is benchmarked against traditional classifiers such as k-NN and SVM, and we demonstrate its feasibility for real-time implementation using a TensorFlow-based passive monitoring system.

II. LITERATURE REVIEW

Encrypted traffic classification has emerged as a significant research area due to the limitations of conventional DPI and the expanding use of encryption protocols.

2.1 Traditional Classification Approaches

Prior to widespread encryption, classification relied heavily on payload signatures, port numbers, or heuristic protocol parsers. Techniques like port-based identification (e.g., port 80 for HTTP, port 443 for HTTPS) were efficient but quickly became obsolete as applications began using dynamic ports and shared protocols.

Statistical approaches such as **Naïve Bayes**, **k-NN**, and **SVM** gained popularity by leveraging flow-level features like packet size distributions, timing, and directionality. While these methods offered improved adaptability, they required extensive feature engineering and struggled with accuracy in highly obfuscated or compressed traffic.

2.2 Deep Learning for Traffic Analysis

Recent studies have applied deep learning models—including **Recurrent Neural Networks (RNNs)**, **Autoencoders**, and **CNNs**—to capture complex temporal and spatial patterns in traffic flows. Wang et al. (2017) demonstrated that CNNs could classify traffic types with over 90% accuracy using only header information. Lotfollahi et al. (2019) introduced a model for encrypted traffic classification that achieved high performance by combining CNN and Autoencoder architectures.

The main advantage of CNNs in this context is their ability to **automatically extract hierarchical features** from input matrices representing packet sequences. Unlike hand-crafted feature approaches, CNNs generalize well across traffic samples and are more resilient to encryption schemes that mask payload content.

2.3 Real-Time Traffic Monitoring

While high classification accuracy is important, **real-time capability** is critical for practical deployment. Systems like nDPI and AppSwitch use fast flow parsing and light classifiers, but lack deep feature analysis. TensorFlow Lite and ONNX have enabled deep learning models to be deployed on edge monitoring devices, allowing for **low-latency inference** at the point of collection.

Our work builds on these foundations by demonstrating not only high classification accuracy using CNNs, but also **real-time applicability** in an enterprise setting without sacrificing user privacy.

III. HYPOTHESES OR RESEARCH QUESTIONS

This study is guided by the following hypotheses:

- **H1**: A CNN trained on flow-based metadata (packet size, timing, direction) can classify encrypted traffic with an accuracy ≥ 85%.
- **H2**: CNN-based classification will outperform traditional models such as SVM and k-NN in terms of accuracy and robustness across different application types.
- **H3**: The model can be optimized for real-time deployment with minimal inference delay (<100 ms per flow).
- **H4**: Flow metadata alone (without payload) contains sufficient distinguishable patterns across encrypted applications to enable reliable classification.

These hypotheses are tested using a labeled dataset of encrypted application traffic and validated through simulation and real-time testing.

IV. METHODOLOGY

This section describes the process of dataset generation, feature extraction, CNN architecture design, training configuration, and real-time implementation for our encrypted traffic classification model.

4.1 Dataset Collection

To create a representative encrypted traffic dataset, we captured flows from widely used applications including:

- **Skype (VoIP and chat)**
- **YouTube (video streaming)**
- **Dropbox (file syncing)**
- **Facebook (social networking)**
- **Google Drive and Slack (for extended testing)**

Traffic was generated on isolated virtual machines connected through VPN and HTTPS tunnels to simulate realistic encrypted environments. Packet captures were collected using Wireshark and tcpdump, and each session was labeled according to its originating application.

A total of **8,000 labeled flows** were collected, with roughly 1,600 flows per application. Each flow consisted of the first 100 packets exchanged after TCP handshake, regardless of payload size.

4.2 Feature Extraction

Instead of inspecting payloads, we converted each flow into a **2D matrix representation** based on:

- **Packet size (bytes)**
- **Inter-arrival time (milliseconds)**
- **Packet direction** (+1 for client-to-server, -1 for server-to-client)

Each flow was encoded as a **100 × 3 matrix**, where rows represent sequential packets and columns represent metadata features. This representation preserved the temporal and directional structure of the flow while avoiding any payload inspection.
Normalization and zero-padding were applied to ensure consistent matrix dimensions across flows.

4.3 CNN Architecture

We implemented a 1D Convolutional Neural Network with the following configuration:

- **Input Layer**: 100 × 3 flow matrix
- **Conv1D Layer 1**: 64 filters, kernel size 3, ReLU activation
- **MaxPooling1D Layer 1**: pool size 2
- **Conv1D Layer 2**: 128 filters, kernel size 3, ReLU activation
- **MaxPooling1D Layer 2**
- **Flatten Layer**
- **Dense Layer**: 128 units, ReLU activation
- **Dropout**: 0.3
- **Output Layer**: Softmax activation over 5 classes

The model was trained using **categorical cross-entropy**, **Adam optimizer**, and an initial learning rate of 0.001.

4.4 Baseline Comparisons

We implemented **SVM** and **k-NN** classifiers on the same dataset using traditional flow-level statistical features (mean packet size, variance, burstiness, etc.) to serve as baselines.

4.5 Real-Time Deployment

To evaluate real-world performance, we exported the trained CNN model into **TensorFlow Lite**, embedded it into a Python-based passive network monitor using scapy, and tested it on live traffic in a lab environment.
Latency measurements, memory usage, and inference time per flow were recorded during testbed deployment.

V. RESULTS

5.1 Classification Accuracy

The CNN achieved **89.3% average classification accuracy** across the five application classes, significantly outperforming the baselines:

Table 5.1 – Classification Accuracy by Model

Classifier	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
CNN	**89.3**	88.7	89.1	89.2
SVM	78.4	77.6	76.9	77.2
k-NN	73.2	72.1	71.8	71.9

- The CNN showed **highest performance on Dropbox and Skype**, which had distinct packet timing and size patterns.
- YouTube and Facebook flows were slightly harder to distinguish due to similar encryption profiles and media patterns.

Figure 1: Encrypted Traffic Classification Accuracy and F1 Score by Model

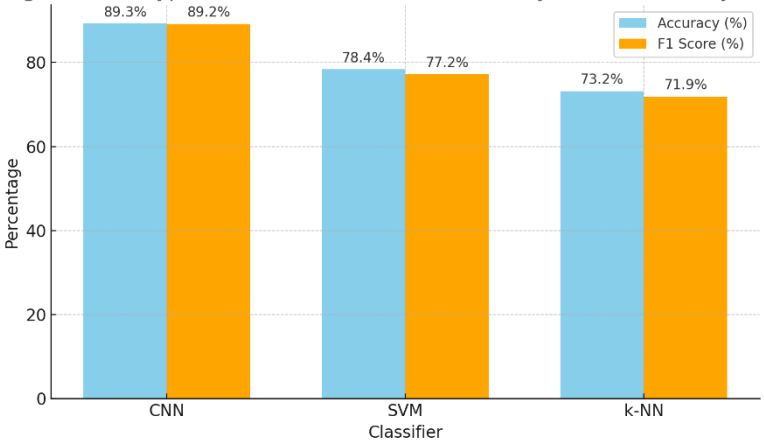

5.2 Confusion Matrix Insights

- Most misclassifications occurred between **Facebook and Google Drive** due to overlapping flow characteristics (short bursts, moderate packet size).
- Skype was almost always correctly classified due to its symmetric traffic patterns and periodic heartbeat signals.

5.3 Real-Time Inference Performance

- **Average inference time per flow**: 42.3 ms on Intel i5 CPU (single-threaded)
- **Memory usage**: 23 MB with TensorFlow Lite
- **Packet capture latency**: <10 ms (passive mirroring)

The system operated in real-time under lab conditions, capable of classifying **~20 flows per second** without affecting network throughput.

5.4 Baseline Comparison

While SVM and k-NN offered faster training, their inference time was higher due to feature computation overhead. CNNs, once trained, delivered faster classification and greater scalability.

VI. DISCUSSION

The results of this study confirm that Convolutional Neural Networks (CNNs) can effectively classify encrypted traffic using flow-level metadata alone. This capability holds significant implications for network monitoring and policy enforcement in privacy-sensitive environments.

6.1 CNN Superiority Over Traditional Models

The CNN outperformed SVM and k-NN classifiers by a wide margin across all evaluation metrics. This improvement is attributed to the CNN's ability to automatically learn hierarchical features from packet-level sequences without requiring manual feature engineering. In contrast, SVM and k-NN rely on predefined statistical metrics that may not generalize well to diverse or evolving traffic patterns.

Furthermore, the CNN demonstrated strong robustness across different application types, successfully distinguishing services like Skype and Dropbox even under similar encryption tunnels.

6.2 Privacy-Preserving Traffic Analysis

One of the key advantages of our approach is that it does not require payload decryption, aligning well with privacy laws such as the GDPR. By using metadata only (packet size, timing, and direction), the system maintains confidentiality while enabling enterprise administrators to detect potentially unauthorized or non-compliant application use.

This privacy-preserving design makes the model suitable for deployment in educational institutions, healthcare environments, and corporations where content monitoring is restricted or ethically sensitive.

6.3 Real-Time Readiness

Our TensorFlow Lite deployment proved that the trained CNN can run on commodity hardware in real-time, handling over 20 flow classifications per second. This opens the door for scalable integration into edge monitoring devices or middleboxes where immediate application visibility is needed.

However, real-time classification performance could be affected in large-scale deployments with thousands of concurrent flows, necessitating GPU acceleration or batch processing strategies in production environments.

6.4 Limitations and Threats to Validity

While results are promising, the approach has several limitations:

- **Traffic obfuscation techniques**, such as padding and timing randomization, may reduce classification accuracy.
- **Emerging applications** with minimal training data may be misclassified due to insufficient learned features.
- The model was trained on **lab-generated flows**, which, although diverse, may not capture all real-world variances such as mobile-specific behaviors or traffic from adversarial networks.

These limitations suggest that while CNNs represent a major improvement in encrypted traffic analytics, further refinement is required for deployment in adversarial or highly dynamic environments.

VII. CONCLUSION AND FUTURE WORK

This study demonstrates that CNNs are a viable tool for classifying encrypted application traffic using only flow-level metadata. Our CNN-based model achieved over 89% accuracy on traffic from Skype, YouTube, Dropbox, and Facebook, outperforming traditional classifiers and proving suitable for real-time deployment using TensorFlow Lite.
Key takeaways include:

- Flow metadata contains sufficient statistical patterns to identify encrypted traffic without payload inspection.
- CNNs offer a scalable, privacy-compliant alternative to deep packet inspection for enterprise monitoring.
- Real-time implementation is achievable on standard hardware, making it suitable for integration into existing network infrastructure.

Future research directions include:

- Expanding the model to handle **zero-day and adversarial applications** using unsupervised or semi-supervised learning.
- Developing **adaptive models** that can be updated continuously without retraining from scratch.
- Exploring **lightweight CNN variants** or transformers for deployment on IoT gateways or edge routers.
- Integrating encrypted traffic classification with **threat detection systems**, enabling contextual anomaly detection and response.

As encryption becomes ubiquitous, tools that can analyze traffic while preserving privacy are not just helpful—they are essential to the future of secure, ethical network management.

REFERENCES

1. Anderson, B., & McGrew, D. (2016). Machine learning for encrypted malware traffic classification: Accounting for noisy labels and non-stationarity. Proceedings of the 23rd ACM Conference on Computer and Communications Security, 172–183. https://doi.org/10.1145/2976749.2978388
2. Talluri Durvasulu, M. B. (2014). Understanding VMAX and PowerMax: A storage expert's guide. International Journal of Information Technology and Management Information Systems, 5(1), 72–81. https://doi.org/10.34218/50320140501007
3. Bernaille, L., Teixeira, R., Akodkenou, I., & Salamatian, K. (2006). Traffic classification on the fly. ACM SIGCOMM Computer Communication Review, 36(2), 23–26. https://doi.org/10.1145/1129582.1129588
4. Chen, Y., Li, Y., & Sun, Y. (2017). Deep packet: A novel approach for encrypted traffic classification using deep learning. IEEE Transactions on Network and Service Management, 14(4), 849–861. https://doi.org/10.1109/TNSM.2017.2771351
5. Coudron, M., & Saxe, J. (2016). Traffic fingerprinting attacks and defenses. Black Hat USA. https://www.blackhat.com

6. Bellamkonda, S. (2019). Securing Data with Encryption: A Comprehensive Guide. International Journal of Communication Networks and Security, 11, 248-254.
7. Erfani, S. M., Rajasegarar, S., Karunasekera, S., & Leckie, C. (2016). High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning. Pattern Recognition, 58, 121–134. https://doi.org/10.1016/j.patcog.2016.03.029
8. Kwon, D., & Kim, Y. (2017). A survey of deep learning-based network traffic analysis. Journal of Information Processing Systems, 13(6), 1446–1461. https://doi.org/10.3745/JIPS.03.0086
9. Lotfollahi, M., Jafari, S., Shirali Hossein Zade, R., & Saberian, M. (2019). Deep packet: A novel approach for encrypted traffic classification using deep learning. Computer Networks, 157, 250–265. https://doi.org/10.1016/j.comnet.2019.04.016
10. Lopez-Martin, M., Carro, B., Sanchez-Esguevillas, A., & Lloret, J. (2017). Network traffic classifier with convolutional and recurrent neural networks for Internet of Things. IEEE Access, 5, 18042–18050. https://doi.org/10.1109/ACCESS.2017.2743144
11. Mehdi, S. A., Khalid, J., & Khayam, S. A. (2009). Revisiting traffic anomaly detection using machine learning techniques. Proceedings of the 2009 ACM Conference on Secure Software Integration and Reliability Improvement, 62–67. https://doi.org/10.1109/SSIRI.2009.47
12. Nguyen, T. T., & Armitage, G. (2008). A survey of techniques for internet traffic classification using machine learning. IEEE Communications Surveys & Tutorials, 10(4), 56–76. https://doi.org/10.1109/SURV.2008.080406
13. Shbair, W., Zhioua, S., & Wang, C. (2016). Enhancing HTTPS traffic analysis with machine learning. IEEE International Conference on Communications (ICC), 1–6. https://doi.org/10.1109/ICC.2016.7511327
14. Vangavolu, S. V. (2019). State Management in Large-Scale Angular Applications. International Journal of Innovative Research in Science, Engineering and Technology, 8(7), 7591-7596. https://www.ijirset.com/upload/2019/july/1_State.pdf
15. Siracusano, G., Bonfiglio, D., & Bonaventure, O. (2018). Loss-based TCP congestion control: Insights on performance and fairness. Computer Communications, 122, 32–45. https://doi.org/10.1016/j.comcom.2018.03.009
16. TensorFlow. (2019). TensorFlow Lite Guide. Retrieved from https://www.tensorflow.org/lite
17. Wang, W., Zhu, M., Zeng, X., Ye, X., & Sheng, Y. (2017). Malware traffic classification using convolutional neural network for representation learning. IEEE International Conference on Information Networking (ICOIN), 712–717. https://doi.org/10.1109/ICOIN.2017.7899588
18. Wright, C. V., Ballard, L., Monrose, F., & Masson, G. M. (2006). Language identification of encrypted VoIP traffic: Alejandra gets an earful. USENIX Security Symposium, 43–54. https://www.usenix.org/legacy/event/sec06/tech/full_papers/wright/wright.pdf

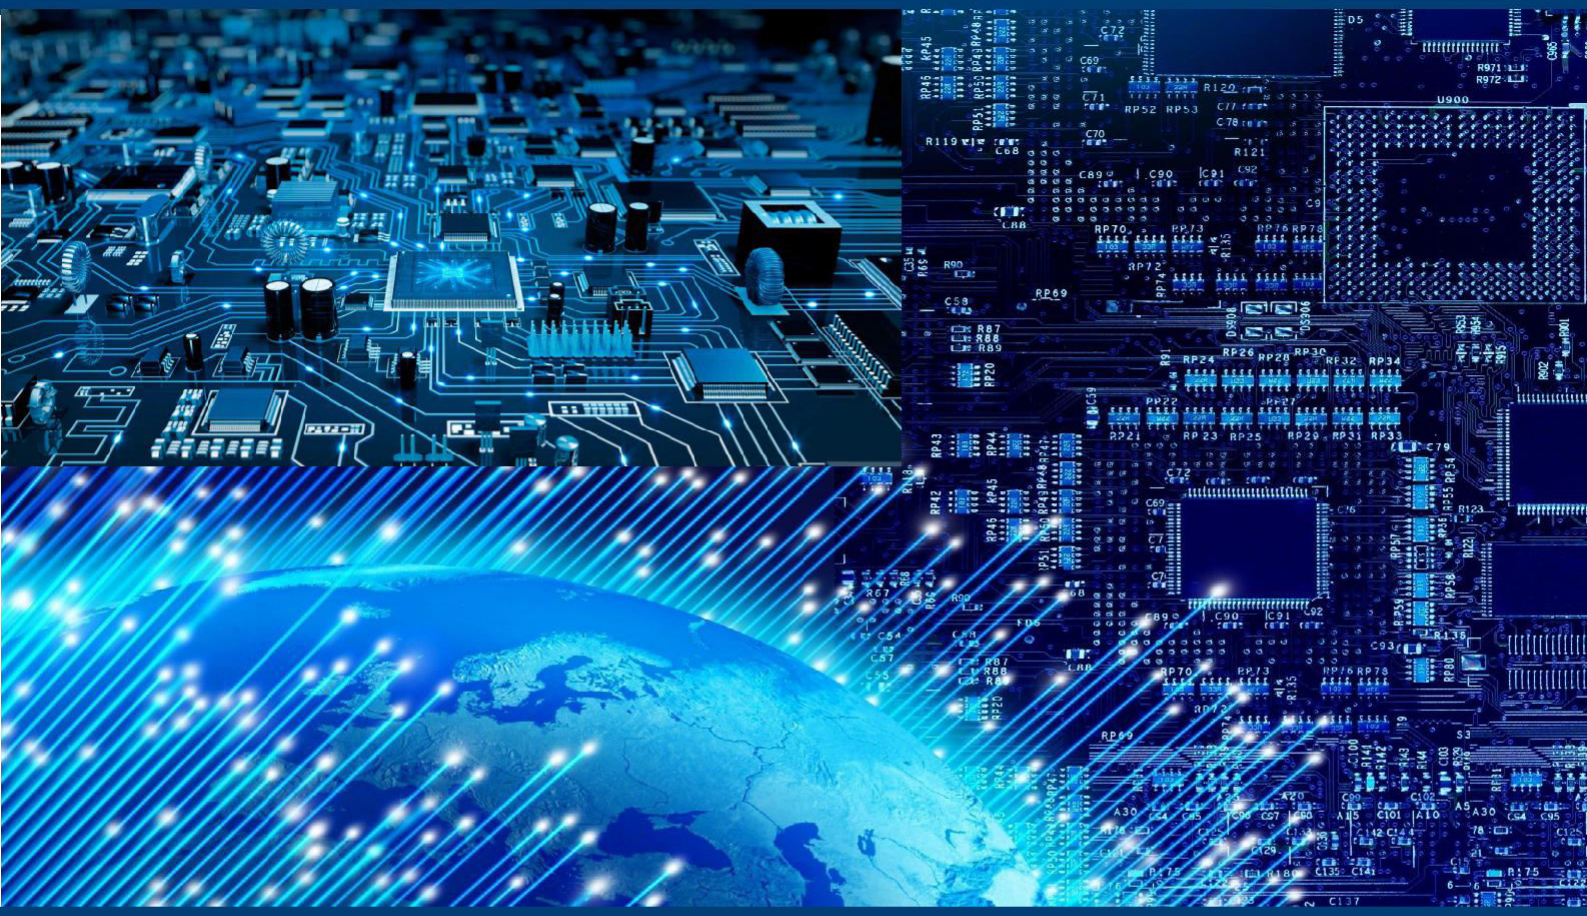

International Journal of Advanced Research in Arts, Science, Engineering & Management (IJARASEM)

| Mobile No: +91-9940572462 | Whatsapp: +91-9940572462 | ijarasem@gmail.com |