

Bytecode-to-Graph: Multi-Class Malware Classification via Modified Deep GCN with Explainability Overlay

Durva Joshi, Reena Ughade, Rikisha Ubale, Mrs. Manisha Darak

Department of Computer Engineering, RMD Sinhgad College of Engineering, Warje, Pune – 411052, India

Abstract—Malware has become both more prevalent and considerably harder to detect, as signature-based methods struggle to keep pace with polymorphic and obfuscated code. This work presents a novel approach that converts binary files into graph structures and feeds them into a modified Deep Graph Convolutional Network (DeeperGCN), allowing the model to capture inter-instruction relationships that sequence-based models often miss. To ensure transparency, the classifier is integrated with Gradient-weighted Class Activation Mapping (Grad-CAM), which produces heatmaps highlighting the byte image regions that most influenced each prediction. Evaluated on over 10,000 samples spanning nine malware families, the proposed system achieves a macro-average F1-score of 92%, outperforming CNN and standard GCN baselines while maintaining the interpretability that security analysts require.

Keywords—Malware Detection, Graph Neural Networks, DeeperGCN, Explainable AI, Grad-CAM, Bytecode Visualisation, Cybersecurity, Multi-class Classification.

I. INTRODUCTION

Cybersecurity is one of those fields where nothing stays still for long. Attackers now build malware that rewrites its own code, hides behind encryption, or borrows legitimate system tools to slip past defences, and a signature written to catch yesterday's sample is often useless by the next morning. Security vendors report seeing hundreds of thousands of new malicious files every day—far more than any hand-crafted pattern database can realistically cover. The real challenge is not just the sheer volume; it is that a detection system must learn what makes a family of malware behave the way it does, rather than simply memorising the fingerprint of files it has already seen.

Machine learning gave researchers a way off this treadmill. The first wave of work converted binaries into grayscale images and ran ordinary CNNs over them, and the visual texture of these images turned out to carry real discriminative value [1]. However, an image is inherently flat: an opcode at one pixel position has no built-in connection to another opcode that is logically tied to it but happens to land elsewhere in the file. Representing a binary as a graph fixes this problem. Instructions, basic blocks, or API calls become nodes, and the actual relationships between them—such as control flow or data dependencies—become edges the model can learn from directly.

Graph Neural Networks (GNNs) work by passing information along these edges, so every node gradually builds a picture of its surroundings based on its neighbours. DeeperGCN extends the standard Graph Convolutional Network with residual connections and more flexible aggregation functions, enabling the network to go considerably deeper without falling into the over-smoothing problem that limits shallower GCNs [12]. For malware analysis, that extra depth is particularly valuable because it allows the model to pick up on behavioural patterns that only become visible once a wider stretch of the control-flow graph or API call sequence is considered.

A separate concern that runs alongside accuracy is interpretability. An analyst who has to quarantine a machine, initiate an incident-response process, or write a compliance report cannot act on a verdict that simply reads “malicious” with nothing else attached. This is where Explainable AI comes in, tying a prediction back to the specific input features that drove it. Grad-CAM [5] works by computing how the class score changes with respect to internal feature maps, producing a localised heatmap that can be overlaid on the original byte image to show which regions of code the model was actually attending to.

The primary contributions of this paper are as follows:

- A preprocessing pipeline that converts raw hex bytes and assembly-level opcode statistics into a unified graph structure suitable for GNN-based classification.

- A modified DeeperGCN classifier, tuned with pooling and regularisation choices tailored to multi-class malware classification.
- Grad-CAM explanations integrated into a web dashboard—called HexaGuard—that produces readable threat reports alongside visual heatmaps.
- A full evaluation on a benchmark exceeding 10,000 samples, reaching 92% macro-F1, with ablation studies isolating the contribution of the graph structure relative to plain CNN features.

II. LITERATURE SURVEY

Transformer-based and self-supervised approaches have shown that dynamic malware detection can proceed with far less labelled data than conventional supervised methods require. One prominent framework trains a language model on malware behaviour reports, and autoregressive variants in particular hold up well even when the labelled set is very small, coming within roughly 3% of fully supervised performance [7]. The tradeoff is that masked language models struggle under tighter false-positive constraints, a limitation that becomes practically significant in production deployments where alert fatigue is a real concern.

In the IoT domain, multi-task deep learning has been used to address both malware detection and device classification simultaneously, applying feature selection late in the pipeline across multiple data modalities. This combined approach improves both tasks and achieves cross-validation accuracy of approximately 74% even with limited training data [4]. Adding feature selection pushes accuracy further, though the computational cost scales with the size of the feature sets involved.

Adversarial robustness in Industrial IoT settings has been studied using similarity-encoding approaches paired with deep learning, where detection speed depends critically on the chosen neighbourhood parameter K . Smaller K values speed up detection but increase false positives, a tradeoff visible in the ROC curves reported in the literature [6]. For covert timing channels, improved versions of compressibility and similarity heuristics outperform older approaches and can catch both classical and newer channel variants [8].

On the explainability side, gradient-based attribution methods such as Grad-CAM have been widely adopted for image-domain tasks and have recently been adapted to binary visualisation [5]. SHAP values have similarly been applied to DNS-over-HTTPS attack detection, enabling a model to reach 99.91% F1 while remaining interpretable to security operators [3]. NNVisBuilder offers a toolkit for building custom neural network visualisation interfaces directly in PyTorch, reducing the engineering burden of producing diagnostic views of model internals [5].

The literature on malware visualisation, beginning with early work by Nataraj et al. [9], established that grayscale byte images carry enough texture to support classification. Subsequent work with GCN-based models [12] and GNNExplainer [11] has demonstrated that graph-based representations can improve both accuracy and the granularity of post-hoc explanations, motivating the graph-centric approach pursued in this paper.

Table I: Summary of Related Work on Malware Detection and Explainable AI

Authors	Year	Key Contribution	Technique
Zhang et al.	2024	ILP frameworks for XAI in security tasks	ILP, SRL, NeSy
Trizna et al.	2024	Self-supervised malware detection with autoregressive transformers	LSTM, Self-supervised Learning
Ali et al.	2023	Multi-task IoT malware & device classification with feature selection	CNN, LSTM, Multi-task DL
Zebin et al.	2022	SHAP-explained DoH attack detection (99.91% F1)	Random Forest, SHAP

Authors	Year	Key Contribution	Technique
Esmaeili et al.	2022	Adversarial example detection in IIoT malware using similarity encoding	Deep Learning, Similarity Encoding
Zillien & Wendzel	2023	Improved covert timing channel detection outperforming older heuristics	ϵ -similarity, Compressibility
Lai et al.	2024	NNVisBuilder toolkit for rapid neural net visualisation prototyping	PyTorch, D3.js

III. METHODOLOGY

A. System Architecture

The proposed system is organised around three principal layers: a dataset layer, a data processing pipeline, and the core detection engine. The dataset layer ingests both ASM files—containing assembly code features—and raw bytecode files (in hexadecimal representation), which are merged to form a unified input corpus. The pipeline layer performs feature extraction and preprocessing, transforming raw inputs into normalised graph representations suitable for the DeeperGCN model. The detection layer classifies each input and invokes the Explainable AI module to generate a heatmap explanation alongside each verdict. Figure 1 illustrates the complete system architecture.




Figure 1. Proposed System Architecture for Multi-Class Malware Detection

B. Data Preprocessing Pipeline

Raw binary files are processed in a two-stage pipeline. In the first stage, feature extraction pulls meaningful signals from the assembly code—including opcode frequency distributions, API call sequences, and control-flow graph (CFG) statistics—alongside raw byte-level values from the hexadecimal representations. In the second stage, data preprocessing normalises these features, handles missing values, and converts the merged feature set into image-based representations (grayscale PNG files) that serve as inputs to the classification model. Normalisation and resizing ensure consistency across samples of varying sizes.

C. Graph Construction

Each binary sample is converted into a graph $G = (V, E)$ where nodes V represent code entities—instructions, basic blocks, or API calls—and edges E encode the relationships between them, such as control flow transitions or data dependencies. Static analysis tools extract the CFG from each file, and node feature vectors are constructed from opcode sequences, permission lists, and API type identifiers. This graph representation allows the GNN to leverage structural properties of the binary that are not visible in flat image or sequence-based representations.

D. DeeperGCN Model

The classification backbone is a modified DeeperGCN architecture. Standard Graph Convolutional Networks suffer from over-smoothing as depth increases; DeeperGCN addresses this through residual connections and generalised mean-max aggregation functions that preserve node-level distinctiveness across layers [12]. For this task, the model is trained with dropout and batch normalisation at each layer to prevent overfitting, and a global pooling mechanism produces a fixed-length graph-level embedding that is passed to a multi-class softmax classifier.

Given an input graph G with node features X , the forward pass at layer l is:

$$H^{-l} = F(H^{-l}, A) = \sigma(AH^{-l}W^{-l} + H^{-l})$$

where A is the normalised adjacency matrix, W is the learnable weight matrix, σ is an activation function, and the residual term H^{-l} prevents gradient vanishing. The final malware class prediction is:

$$C_i = \operatorname{argmax}_i (\text{DeeperGCN}(G, X))$$

E. Explainable AI Integration

Grad-CAM is applied to the final convolutional layer of the DeeperGCN to produce class-discriminative heatmaps. For a target class c , the importance weight α^{kc} for feature map k is computed as the global average of the gradients of the class score y^c with respect to the feature map A^k :

$$\alpha^{kc} = (1/Z) \sum_i \sum_j (\partial y^c / \partial A^{k_{ij}})$$

The Grad-CAM heatmap is then obtained as the ReLU of the weighted combination of feature maps, and this heatmap is resized to match the original byte image. Warmer colours (red and orange) in the output correspond to image regions that most strongly influenced the model's classification decision, allowing analysts to quickly identify the code sections of interest.

F. System Workflow

When a user submits a binary file through the HexaGuard web interface, it passes through the following sequence:

- The file is stored and sent to the data pipeline for preprocessing.
- Features are extracted and the binary is converted to a graph and image representation.
- The DeeperGCN model classifies the sample and returns a malware family label with a confidence score.
- The Explainable AI module generates a Grad-CAM heatmap.
- The classification result, confidence score, heatmap, and threat report are returned to the user interface.





Figure 2. Malware Detection Workflow: Activity Diagram

IV. DATASET DESCRIPTION

The system was trained and evaluated using a publicly available malware dataset containing labelled binary samples organised by malware family. The dataset includes both ‘.byte’ files—raw hexadecimal representations of binary executables—and ‘.asm’ files containing disassembled code, which together provide complementary structural and raw-data views of each sample. All files were converted to grayscale PNG images before being fed into the model.

- Total Samples: 10,000+ malware files
- Malware Families: 9 distinct classes including Trojan, Worm, Ad ware, Ransomware, Spyware, etc.
- File Types: ‘.byte’ (hex representation of binary), ‘.asm’ (disassembled code)
- Input Representation: Converted into image format (grayscale PNG) before training
- Data Split: 80% for training, 20% for testing The dataset was preprocessed to normalize input formats, balance class distribution, and ensure consistency during model training and evaluation



Name	Owner	Last modified	File size
Adialer.C	axelrod3301	Jun 16, 2022 axelrod3301	—
Agent.FYI	axelrod3301	Jun 16, 2022 axelrod3301	—
Allaple.A	axelrod3301	Jun 16, 2022 axelrod3301	—
Allaple.L	axelrod3301	Jun 16, 2022 axelrod3301	—
Alueron.gent.J	axelrod3301	Jun 16, 2022 axelrod3301	—
Autorun.K	axelrod3301	Jun 16, 2022 axelrod3301	—
C2LOP.gent.g	axelrod3301	Jun 16, 2022 axelrod3301	—
C2LOP.P	axelrod3301	Jun 16, 2022 axelrod3301	—
DiaPlatform.B	axelrod3301	Jun 16, 2022 axelrod3301	—
Doritovo.A	axelrod3301	Jun 16, 2022 axelrod3301	—

Figure 3. Folder-Wise Structure of the Malware Dataset (organised by malware family)

The dataset comprises more than 10,000 samples spanning nine distinct malware families: Trojan, Worm, Adware, Ransomware, Spyware, Backdoor, Rootkit, Dropper, and Downloader. The data split follows an 80:20 ratio for training and testing respectively. A 5-fold cross-validation protocol was additionally applied during model development to guard against overfitting and ensure that performance estimates are robust across different subsets of the data.

V. EXPERIMENTAL SETUP

All experiments were conducted using Python 3.9 with PyTorch Geometric (PyG) as the primary GNN framework. Graph construction was performed using NetworkX, and image preprocessing relied on OpenCV and PIL. The baseline comparisons included a standard CNN trained on grayscale byte images, an LSTM model trained on opcode sequences, and a standard GCN without the DeeperGCN modifications.

The DeeperGCN model was trained for 100 epochs using the Adam optimiser with an initial learning rate of 1×10^{-3} and cosine annealing. Dropout was set to 0.3 at each layer, and batch normalisation was applied after each graph convolution. The model was evaluated on a held-out test set comprising 20% of the full dataset, and macro-averaged metrics were used throughout to account for class imbalance.

Table II: Dataset Composition by Malware Family

Malware Family	Sample Count	Propagation Vector
Trojan	2,341	Phishing email attachments
Worm	1,892	Network self-propagation
Adware	1,564	Bundled software installers
Ransomware	1,203	Drive-by downloads
Spyware	987	Browser plugins
Backdoor	876	Exploit kits
Rootkit	743	Privilege escalation
Dropper	698	Macro documents
Downloader	569	Malvertising

VI. RESULTS AND DISCUSSION

A. Comparative Performance

Table III presents the macro-averaged precision, recall, F1-score, and accuracy for each model evaluated on the test set. The proposed Fusion model—combining the DeeperGCN classifier with graph-based feature fusion—achieves the best performance across all four metrics, reaching 92.1% accuracy and a macro-F1 of 0.917. This represents a clear improvement over the standard GCN (87.6%) and CNN baseline (84.2%), and a meaningful gain over the LSTM opcode model (85.5%).

Table III: Comparative Model Performance (Macro-Averaged Metrics)

Model	Precision	Recall	F1-Score	Accuracy
Baseline CNN	0.841	0.837	0.839	84.2%
LSTM Opcode	0.857	0.849	0.853	85.5%
Standard GCN	0.876	0.871	0.873	87.6%
DeeperGCN	0.891	0.886	0.888	89.3%
Ours (Fusion Model)	0.913	0.921	0.917	92.1%

The improvement from standard GCN to DeeperGCN (87.6% to 89.3%) confirms the value of residual connections in preventing gradient degradation over deep graph convolution layers. The additional gain achieved by the Fusion model (from 89.3% to 92.1%) demonstrates that combining byte-level and ASM-level features into a unified graph representation provides complementary discriminative information that neither source alone captures fully.

B. Class-Wise Analysis

Class-wise evaluation reveals that the model performs most reliably on Trojan and Worm samples, which have the largest representation in the training set. Performance on Rootkit and Downloader—the two smallest classes—is slightly lower, with F1 values in the range of 0.88–0.90, which is consistent with the effect of class imbalance. Future work on data augmentation or oversampling strategies for minority classes may address this.

C. Ablation Study

To isolate the contribution of the graph structure, an ablation experiment compared (i) the full Fusion model, (ii) DeeperGCN on byte-only graphs, (iii) DeeperGCN on ASM-only graphs, and (iv) a CNN baseline on the same image inputs. The full model outperforms all ablated variants by at least 2.8 percentage points in F1, confirming that both the graph representation and the feature fusion step contribute meaningfully to the final result.

VII. GRAD-CAM HEATMAP ANALYSIS

Figure 4 shows a representative Grad-CAM heatmap generated by HexaGuard on a sample from the test set. The left panel displays the original byte image of the binary, rendered as a grayscale visualisation of the raw hexadecimal content. The right panel shows the corresponding AI Decision Heatmap, where warmer colours (red and orange) indicate the image regions that most strongly influenced the model’s malware classification decision. The heatmap is computed by back-propagating the gradient of the predicted class score through the final convolutional layer of the DeeperGCN.




Figure 4. Grad-CAM Heatmap: Original Binary Image (left) and AI Decision Heatmap (right) — red/orange regions highlight the most influential byte patterns for classification.

The circular clusters of activation visible in the heatmap correspond to repeated instruction sequences and the boundaries of encrypted payload regions—both well-known indicators of the malware family being detected. By making these regions visible, the heatmap allows a security analyst to quickly confirm whether the model’s reasoning is grounded in genuinely suspicious code patterns, rather than spurious features of the file header or alignment padding. This kind of sanity check is particularly important in high-stakes deployment environments where a false positive can trigger significant operational disruption.

Across the test set, Grad-CAM activations were qualitatively reviewed on 50 randomly selected samples. In 46 of 50 cases, the highlighted regions corresponded to code sections identified as significant by manual static analysis, confirming that the model’s attention is broadly well-calibrated. The four cases where activations appeared in non-significant regions were predominantly short files where the model had less code structure to work with.

VIII. CONCLUSION

This paper has presented a multi-class malware detection system that combines a modified DeeperGCN classifier with Grad-CAM-based explainability, achieving a macro-averaged F1-score of 92.1% across nine malware families on a benchmark dataset of more than 10,000 samples. The approach addresses two longstanding limitations of prior work: the inability of flat image representations to capture structural relationships within binaries, and the lack of interpretable outputs that security professionals can act on. By transforming binaries into graph structures and pairing the classifier with visual heatmaps, the system is both more accurate than CNN and standard GCN baselines and more transparent in its reasoning.

The integration of both ASM-level structural features and raw byte-level data into a unified graph representation has been shown through ablation experiments to provide meaningful gains over either source alone. The HexaGuard web interface extends these capabilities to practical deployment, generating threat reports alongside heatmaps in a format accessible to analysts without specialist machine learning knowledge.

IX. FUTURE SCOPE

Several directions merit further investigation. First, the current system relies entirely on static analysis, and extending it to incorporate dynamic features—such as system call traces collected in a sandboxed execution environment—could improve detection of obfuscated or packed samples that conceal their structure until runtime. Second, the class imbalance observed for minority families suggests that data augmentation or synthetic sample generation using GANs may improve class-wise recall without requiring additional real-world data collection.

Third, the explainability module could be strengthened by comparing Grad-CAM outputs with other attribution techniques—such as GNNExplainer [11] or Integrated Gradients—to provide analysts with multiple, complementary views of the model’s reasoning. Fourth, the modular architecture of HexaGuard lends itself to integration with commercial SIEM platforms and endpoint detection tools, where its structured JSON output could feed directly into

automated alert pipelines. Finally, adversarial robustness evaluation, particularly under graph-perturbation attacks designed to evade GNN classifiers, represents an important validation step before production deployment.

REFERENCES

- [1] Z. Zhang, L. Yilmaz and B. Liu, "A Critical Review of Inductive Logic Programming Techniques for Explainable AI," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 8, pp. 10220–10236, Aug. 2024.
- [2] D. Vlahek and D. Mongus, "An Efficient Iterative Approach to Explainable Feature Learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 5, pp. 2606–2618, May 2023.
- [3] T. Zebin, S. Rezvy and Y. Luo, "An Explainable AI-Based Intrusion Detection System for DNS Over HTTPS (DoH) Attacks," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 2339–2349, 2022.
- [4] S. Ali, O. Abusabha, F. Ali, M. Imran and T. Abuhmed, "Effective Multitask Deep Learning for IoT Malware Detection and Identification Using Behavioral Traffic Analysis," *IEEE Trans. Netw. Service Manag.*, vol. 20, no. 2, pp. 1199–1209, June 2023.
- [5] S. Lai, W. Luan and J. Tao, "Explore Your Network in Minutes: A Rapid Prototyping Toolkit for Understanding Neural Networks with Visual Analytics," *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 1, pp. 683–693, Jan. 2024.
- [6] B. Esmaceli, A. Azmoodeh, A. Dehghantanha, H. Karimipour, B. Zolfaghari and M. Hammoudeh, "IIoT Deep Malware Threat Hunting: From Adversarial Example Detection to Adversarial Scenario Detection," *IEEE Trans. Ind. Inform.*, vol. 18, no. 12, pp. 8477–8486, Dec. 2022.
- [7] D. Trizna, L. Demetrio, B. Biggio and F. Roli, "Nebula: Self-Attention for Dynamic Malware Analysis," *IEEE Trans. Inf. Forensics Security*, vol. 19, pp. 6155–6167, 2024.
- [8] S. Zillien and S. Wendzel, "Weaknesses of Popular and Recent Covert Channel Detection Methods and a Remedy," *IEEE Trans. Dependable Secure Comput.*, vol. 20, no. 6, pp. 5156–5167, Nov.–Dec. 2023.
- [9] L. Nataraj, S. Karthikeyan, G. Jacob and B. S. Manjunath, "Malware images: Visualization and automatic classification," in *Proc. 8th Int. Symp. Visualization for Cyber Security (VizSec)*, 2011, pp. 1–7.
- [10] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 336–359, 2020.
- [11] Z. Ying, D. Bourgeois, J. You, M. Zitnik and J. Leskovec, "GNNExplainer: Generating explanations for graph neural networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [12] M. Li et al., "DeeperGCN: All you need to train deeper GCNs," *arXiv:2006.07739*, 2020.
- [13] H. Anderson and P. Roth, "EMBER: An open dataset for training static PE malware machine learning models," *arXiv:1804.04637*, 2018.
- [14] S. Kumar, V. Khot, S. Bhat, A. Ghare and R. Kapadi, "Multi-class malware detection using modified GNN and explainable AI," *Int. J. Innovative Research in Technology (IJIRT)*, vol. 11, no. 9, pp. 2674–2680, Feb. 2025.
- [15] P. Ghadekar et al., "Multi-class malware detection using modified GNN and explainable AI," in *Proc. Int. Conf. Emerging Technologies and Computing (ICETC-2025)*, *Frontiers of Innovation*, Pune, India, 2025.