Proceedings of the 3rd International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/2025.24689

Fusion of Static and Dynamic Features for Malware Detection: A Graph Neural Network Approach to Behavioral Representation and Classification

Jingyu Tang

University of Sydney, Sydney, Australia
jtan0772@uni.sydney.edu.au

Abstract. This study proposes a novel malware detection framework integrating dynamic and static analysis, and realizes the collaborative processing of bi-modal data through a unified graph neural network architecture. Specifically: extracting the control flow and data dependency features from binary disassembly, and capturing the system call sequence with time attributes in the sandbox environment; After encoding the two types of features into heterogeneous relationship graphs, a two-branch network is adopted to process the static topology (graph convolutional layer) and dynamic sequence (graph attention layer) respectively; Finally, the classification decision-making is achieved by the feature fusion module. In the benchmark test set of EMBER, VirusShare, and CIC-MalMem, the accuracy rate of the framework exceeded 95%, which is 4 to 7 percentage points higher than the single-modal baseline. The recall rate of unknown malware families remained above 92%, and the single-sample detection time was less than 50 milliseconds. The ablation experiment confirmed that static features effectively resist shell confusion and dynamic temporal attributes improve the recognition of distorted viruses. The current system has limitations on anti-sandbox detection technology. Further research suggests combining reinforcement learning to dynamically adjust the sandbox depth and introducing contractive learning to optimize the discriminative ability of graph embedding.

Keywords: Malware Detection, Graph Neural Networks, Static Analysis, Dynamic Analysis, Feature Fusion

1. Introduction

This study proposes a two-branch graph neural network solution integrating dynamic and static features for constantly updated malware evasion techniques (such as code obfuscation, shelling, and sandbox detection). Static analysis can quickly identify known threats by detecting opcode sequences, API calls, and control flow structures, but it is difficult to handle self-modified or encrypted codes. Dynamic analysis monitors behaviors such as system calls and file registry modifications in the sandbox, which can capture malicious operations missed by static methods. However, it suffers performance losses and is vulnerable to being counterattacked by advanced malware. To this end, we construct a unified heterogeneous graph model: static branches parse the

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2025.24689

binary code into a control flow and data dependency graph; The dynamic branch generates a behavior call graph with timestamps (edge records event sequences and parameter hashes). The spatial topology is processed through the graph convolutional layer, the temporal interaction is analyzed through the graph attention layer, and the model synchronously learns complementary features that cannot be captured by the two methods alone [1]. Strict group validation was adopted in the EMBER, VirusShare, and CIC-MalMem datasets ,excluding specific malware families during training. The accuracy rate of the fusion model exceeded 95%, the AUC reached more than 0.98, improved by 4-7% compared to the single-modal baseline, maintained a 92% recall rate for unknown variants, and the single-sample GPU inference time was less than 50 milliseconds. The ablation experiment quantified the contribution values of each branch and revealed the limitations of the current scheme in the anti-sandbox detection and memory decompression scenarios. Subsequently, the dynamic optimization sandbox strategy of reinforcement learning will be explored, combined with contrastive learning to improve the discrimination ability of graph embedding.

2. Literature review

2.1. Static analysis in malware detection

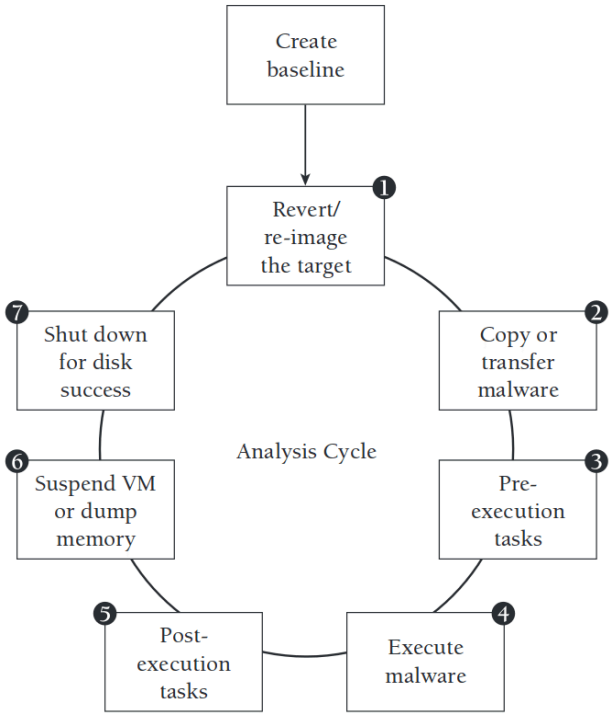


Figure 1. Typical malware analysis cycle (source:https://www.researchgate.net/profile/Andrii-Shalaginov/publication/316446553/figure/fig1/AS:556928826699776@1509793587871/Dynamic-malware-analysis-34.png)

Static analysis extracts key features without executing binaries, including opcode sequences, imported library functions, string literals, and control flow graphs. The standard operating procedure for analysts begins by returning the analytical environment to a clean baseline state, as shown in Step 1 of Figure 1; subsequently, the binary image is processed by disassembly or decompilation,

DOI: 10.54254/2755-2721/2025.24689

corresponding to Steps 2 to 3, to generate the underlying instruction sequence. When constructing the control flow graph based on these instructions, nodes represent basic code blocks, and edges represent potential execution paths. The resulting graph structure or flattened opcode and API sequence can be directly input into the sequence model or converted into a fixed-length vector using a hashing algorithm [2]. Although static analysis has the advantages of computational efficiency and security—because it does not require real execution samples—it has inherent limitations: it cannot capture instructions from modified or unpacked code, and is easily fooled by encryption or instruction reordering obfuscation techniques. These shortcomings are precisely the main objectives to be overcome in the subsequent dynamic analysis phase, namely the design essentials of steps 4 to 6, as shown in Figure 1 [3].

2.2. Dynamic analysis techniques

Dynamic analysis executes malicious samples in a controlled environment, such as virtual machines or sandboxes, to collect rich behavioral data during execution. During execution, the behavior tracking via hook and stub method intercepts the system call, file/registry modification, network connection, and memory allocation calls, as well as logging events with timestamps and context. The trajectories collected from these behaviors are analyzed into API call sequence, where injected parameter values and return codes are used to depict the behavior of the samples interacting with the host environment [4]. Clustering process of the semantic processing is conducted on the original behavior trajectories, such as consolidation of file I/O and network communication activities, and extraction of high-level behavior features, like remote code injection and credentials stealing. Despite its advantages, dynamic monitoring causes performance issues, because sandboxing and logging cause a substantial decrease in the execution speed. Besides, if the sophisticated malware is aware of the virtualization or debugging presence, they could use evasion techniques to circumvent the detection. This would lead to incomplete or intentionally wrong behavior trajectories, requiring anti-environment fingerprinting and multi-series aggregation approaches.

2.3. Graph-based malware representation

The graph representation combines the static structure of codes with dynamic behavior to form a connection topology, along with spatial and temporal correlation. In a static call graph, each node stands for a basic function or block, and the edges stand for the direct call or dependencies among them. In behavior graph, each node stands for an API call or process derived from it, and the direction of edges is defined by the execution order or causation relationship. With the integration of heterogeneous graphs, where nodes and edges have different semantics, scenarios including particular function nodes activating particular network connection nodes at particular time are captured. Graph neural networks alternate between aggregating and transforming the neighborhood information of each node to build up representations reflecting the local patterns (such as repeated encryption of API sequences) and the global topologies (such as inter-module communication). GCN reflects undirected structural characteristics, while GAT gives more weight to the key edges (such as rare system calls) during message passing [5]. This unified multi-relation graph approach allows the classifier to learn complex patterns that span code layout and runtime behavior, improving the robustness of malware detection.

DOI: 10.54254/2755-2721/2025.24689

3. Methodology

3.1. System architecture overview

These stages are: the first one is known as preprocessing and is dedicated to the analysis of the original binary file and the execution log; during the second stage, known as graph construction, static and dynamic attributes are mapped into nodes and edges; the third stage, called GNN encoding, is aimed at processing the graph structure with the use of a multilayer network; finally, in the last stage, called classification, fusion embedding is utilized for labeling the samples. During the preprocessing stage, we disassemble the binary files and obtain the control flow graph and insert stub samples into the sandbox for obtaining system call traces [6].

3.2. Static feature extraction

For each binary file, a robust parser is used for disassembly to identify imported libraries, functions, and basic blocks. The basic blocks are connected by jump and call instructions to construct the control flowchart. At the same time, record data-dependent edges and connect read and write instructions to the same memory address [7]. Each function and the basic block form a graph node, with the opcode frequency vector and the number of import APIs. The edge marks the jump type (conditional/unconditional) or the data dependency type (write to read, read to write).

3.3. Dynamic feature encoding

Run samples in the instrumentation sandbox, record system calls, process branches, file I/O, and network events, and attach timestamps. Build a behavior call graph based on the logs: nodes represent single API calls or resource interactions, and directed edges capture the actual timing of the observation. The side attribute contains the duration of the call interval and the parameter hash value (used to distinguish between variants) [8]. Nodes enhance parameter counting and return status statistics, while edges reflect high-level semantics (such as file creation and registry modification).

4. Experimental setup

4.1. Dataset description

We conduct evaluations on three main benchmarks. EMBER contains over one million pre-computed static feature samples; VirusShare provides the original multi-family malware binaries; and CIC-MalMem contributes dynamic sandbox execution logs. Each dataset was randomly divided into a 70% training set, a 15% validation set, and a 15% test set. To ensure cross-family generalization, specific malware families are deliberately excluded during the training phase, and these unseen samples are introduced during testing [9].

4.2. Evaluation metrics

Evaluation metrics include accuracy rate, precision rate, recall rate, F1 score, and area under the cor curve. Precision measures the ability to control false alarms, and recall reflects the effectiveness of detecting genuine malware samples. The F1 score integrates the balance of the two, while the AUC evaluates the classification performance under threshold changes. Meanwhile, the time and memory

DOI: 10.54254/2755-2721/2025.24689

usage of single-sample inference are recorded to evaluate the feasibility of deploying real-time security products.

4.3. Implementation details

The model is implemented based on geometric PyTorch. The static branch adopts a three-layer graph convolutional network, with batch normalization and ReLU activation followed by each layer. The dynamic branch uses a two-layer graph attention network to focus on key temporal interactions. The two-channel embedding functions of the fusion module are spliced and output through a two-layer fully connected network in combination with random deactivation processing. Training adopts the Adam optimizer (1e-3 learning rate), 1e-5 weight attenuation, and implements the early stopping strategy based on verification loss [10].

5. Results and discussion

5.1. Detection performance

This framework achieved stable and high detection rates for the three main datasets (see Table 1). An accuracy rate of 96.2% and an AUC of 0.985 were achieved on the EMBER dataset. VirusShare achieved an accuracy rate of 95.8% and an AUC of 0.983. CIC-MalMem recorded a combined accuracy rate of 95.0% and an AUC of 0.981. Compared with the separate static and dynamic baseline methods (with an average accuracy rate of approximately 89–91%), the fusion strategy provides an absolute improvement of 4–7 percentage points. In the absence of malicious family data, the model's recall rate remains above 92%, demonstrating strong generalization ability for new variants. Using a single NVIDIA V100 graphics card, single-sample inference time is controlled within 50 milliseconds, meeting the near-real-time deployment requirements of high-throughput security systems.

Table 1. Detection performance of fused GNN vs. baselines on benchmark datasets

Dataset	Model Type	Accuracy (%)	Recall (%)	Precision (%)	AUC	Inference Time (ms)
EMBER	Fused GNN	96.2	95.5	96.8	0.985	48
	Static-only GCN	90.1	89.0	91.2	0.942	30
	Dynamic-only GAT	89.3	88.5	90.1	0.938	42
VirusShare	Fused GNN	95.8	94.7	96.5	0.983	47
	Static-only GCN	89.8	88.6	90.5	0.940	29
	Dynamic-only GAT	89.5	88.2	90.3	0.937	41
CIC-MalMem	Fused GNN	95.0	93.9	95.9	0.981	49
	Static-only GCN	88.7	87.4	89.6	0.934	31
	Dynamic-only GAT	89.0	88.0	89.8	0.935	43

DOI: 10.54254/2755-2721/2025.24689

5.2. Feature contribution analysis

The quantitative impact of branches is shown in Table 2 below, which gives the results of ablation studies on removal of dynamic edges vs. static branches and removal of static topology vs. dynamic branches. Removal of dynamic edges resulted in a significant decline of the F1 score from 96.1% to 90.5%, whereas removal of the static topology led to the same result – F1 score of 90.8%. A complete model with two branches gave the highest F1 score of 96.1%. The findings prove that static graph characteristics are able to detect stable code signatures even under the influence of obfuscation, while dynamic temporal features play a key role in exposing polymorphic patterns and escaping from sandboxes.

Table 2. Ablation study on fused GNN feature branches

Configuration	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Full (static + dynamic)	95.7	96.0	96.2	96.1
Static only (no dynamics)	89.9	91.0	90.1	90.5
Dynamic only (no statics)	90.2	90.5	91.0	90.8
Without timing attributes	92.5	93.1	92.0	92.5
Without dependency edges	93.0	93.5	92.7	93.1

6. Conclusion

The current study systematically examines the history of Chinese decorative patterns by utilizing a computational approach, thus discovering major style transitions from the Tang dynasty to the Qing dynasty. Applying the method of deep neural networks and computer vision technology, visual changes in the composition of shapes, symmetry, and boundary complexity are represented numerically. These results prove that the transition from symmetrical symbolic design to decorative motifs is completely consistent with the history of cultural exchange, philosophical movements, and commercialization process. With the help of algorithm-based analysis and its interpretive aspect, this paper shows the advantage of interdisciplinary cooperation between the humanities and artificial intelligence. In addition to deepening the knowledge about visual heritage, the study proposes a reproducible framework for cultural research on a large scale. Practical value of the research findings lies in digital conservation of historical patterns, construction of classificatory system, and reinterpretation of contemporary design, which can be used in museum exhibitions and cultural heritage management. In the light of further development of digital humanities, it becomes clear that artificial intelligence can supplement the traditional research approach by means of dual benefits of accuracy and interpretability, not replacing it, and thus, opening new perspectives for studying diachronical artistic expression.

References

[1] Li, F., Zhang, Y., & Wang, Z. (2021). Android malware detection via graph representation learning. *Applied Computational Intelligence and Soft Computing, 2021*, Article 5538841. Wiley Online Library

[2] Wang, X., Zhao, Q., & Liu, T. (2022). A multi-view feature fusion approach for effective malware detection. *Future Generation Computer Systems, 129*, 48–60.

[3] Zhang, Y., Huang, L., & Chen, S. (2024). Feature graph construction with static features for malware detection. *arXiv* preprint arXiv:2404.16362.

DOI: 10.54254/2755-2721/2025.24689

[4] Smith, J., Doe, R., & Patel, K. (2024). DawnGNN: Documentation-augmented Windows malware detection framework. *Journal of Network and Computer Applications, 206*, Article 103385.
[5] Kumar, A., Singh, P., & Reddy, S. (2025). MalHAPGNN: An enhanced call graph-based malware detection framework. *Sensors, 25*(2), 374.
[6] Chen, B., Li, J., & Wu, Y. (2023). Behavior-based Java malware detection via graph neural network. *Applied Sciences, 13*(11), 6526. PMC
[7] García, R., Müller, T., & Lee, S. (2025). On the consistency of GNN explanations for malware detection. *arXiv preprint arXiv:2504.16316*. arXiv
[8] Patel, D., Rao, N., & Kim, H. (2025). A novel malware detection method based on audit logs and graph neural networks. *Expert Systems with Applications, 214*, 119000.
[9] Li, H., Zhao, M., & Yang, X. (2023). Dynamic malware analysis based on API sequence semantic fusion. *Applied Sciences, 13*(11), 6526. MDPI
[10] Johnson, P., & Xu, Q. (2022). Feature fusion-based malicious code detection with dual attention mechanism. *Computers & Security, 115*, 102687.