## LAMD: Context-driven Android Malware Detection and Classification with LLMs

Xingzhi Qian<sup>\*†</sup>, Xinran Zheng<sup>\*†</sup>, Yiling He<sup>⊠†</sup>, Shuo Yang<sup>‡</sup> and Lorenzo Cavallaro<sup>†</sup>

<sup>†</sup>University College London <sup>‡</sup>University of Hong Kong <sup>†</sup>{xingzhi.qian.23, xinran.zheng.23, yiling-he, l.cavallaro}@ucl.ac.uk <sup>‡</sup>shuoyang.ee@gmail.com

Abstract—The rapid growth of mobile applications has escalated Android malware threats. Although there are numerous detection methods, they often struggle with evolving attacks, dataset biases, and limited explainability. Large Language Models (LLMs) offer a promising alternative with their zeroshot inference and reasoning capabilities. However, applying LLMs to Android malware detection presents two key challenges: (1) the extensive support code in Android applications, often spanning thousands of classes, exceeds LLMs' context limits and obscures malicious behavior within benign functionality; (2) the structural complexity and interdependencies of Android applications surpass LLMs' sequencebased reasoning, fragmenting code analysis and hindering malicious intent inference. To address these challenges, we propose LAMD, a practical context-driven framework to enable LLM-based Android malware detection. LAMD integrates key context extraction to isolate security-critical code regions and construct program structures, then applies tier-wise code reasoning to analyze application behavior progressively, from low-level instructions to high-level semantics, providing final prediction and explanation. A well-designed factual consistency verification mechanism is equipped to mitigate LLM hallucinations from the first tier. Evaluation in real-world settings demonstrates LAMD's effectiveness over conventional detectors, establishing a feasible basis for LLM-driven malware analysis in dynamic threat landscapes.

#### 1. Introduction

The rapid expansion of the Android ecosystem has heightened security risks, with malware posing serious threats to user privacy, financial security, and sensitive data. Over the past decade, researchers have developed various Android malware detection techniques, yet these methods face persistent challenges in real-world scenarios. Firstly, the open and evolving nature of Android complicates the detection of adaptive malware [1], [2]. Reliance on specific datasets introduces biases, such as ambiguous timestamps and randomly selected samples [3], which further compromise model reliability. Additionally, conventional detectors often lack explainability that fail to offer clear, humanreadable insights into malicious behaviors.

\*Equal contribution and co-first authors. Scorresponding Author.



Figure 1: Two failure cases of applying LLMs to Android malware detection: (1) context window limitations and (2) failure to capture malicious intent. LAMD addresses these challenges by extracting key contexts and tiered reasoning to capture structures and semantics efficiently.

Large Language Models (LLMs) offer a promising paradigm shift in malware detection, differing fundamentally from conventional detectors. They achieve zero-shot inference relying on vast pre-trained knowledge instead of specifically labelled datasets [4], [5], allowing them to handle the evolving malware and potential training bias. Furthermore, to bridge the gap of explainability, the advanced generative capabilities of LLMs present an opportunity by providing human-readable comprehension, thereby enhancing malware analysis from both an accuracy and interpretability perspective.

However, despite the potential, LLMs are not omnipotent. Two primary challenges hinder their effectiveness in Android malware detection: (1) *Excessive support codes in Android applications*: Android malware often comprises thousands of classes to support diverse functionalities across various devices. While some LLMs support up to 2 million tokens [6], directly processing decompiled malware remains impractical, as some samples would still exceed this limit. Truncation offers a potential alternative, but it can cause substantial contextual loss, ultimately degrading analysis accuracy [7]. More importantly, these normal functional codes make malicious codes sparse within the program and become extremely hard to detect. (2) Complex program structures: Although code exhibits structural characteristics akin to natural language [8], it remains fundamentally distinct due to inherent structural complexity [9]. For instance, deeply nested dependencies, intricate API interactions, and class hierarchies all extend beyond sequential token-based modeling. The complexity is further amplified in Android malware, where obfuscation techniques, multi-component interactions, and convoluted function invocations obscure malicious intent. At its core, these challenges raise a fundamental research question:

#### Can we extract crucial semantic and structural information from complete application to guide LLMs in detecting Android malware?

To address this, we draw inspiration from how analysts examine Android malware in the real world, where they identify suspicious APIs, interfaces, and function calls, analyzing contextual relationships to detect malicious behavior within extensive application code. Our framework aims to guide LLMs in replicating this analytical process, enhancing automated explainable Android malware detection.

We propose LAMD, a novel and practical framework that enables LLMs for explainable Android malware detection. LAMD consists of two core components: key context extraction and tier-wise code reasoning. Specifically, we perform static analysis on APKs and employ a custom backward slicing algorithm to extract key variables, dependencies, and invocations for predefined suspicious APIs. These elements are transformed into graphical representations, filtering irrelevant code and preserving essential semantics for the LLM. Furthermore, we introduce tier-wise code reasoning combined with factual consistency verification, transferring structured knowledge across multiple tiers to understand, draw conclusions and explain detection results. This approach refines the LLM's understanding of applications, progressing from fine-grained analysis to higherlevel abstraction. As shown in Figure 1, LAMD effectively addresses real-world challenges, enabling practical and explainable LLM-powered Android malware detection. The main contributions of this paper are as follows:

- We introduce LAMD, the first LLM-powered practical Android malware detection framework, unlocking LLMs' ability for explainable Android malware detection in dynamic scenarios, providing heuristics for LLM-powered malware-related tasks.
- LAMD integrates key context extraction and tierwise code reasoning to filter irrelevant functionalities while capturing semantics and structural dependen-

cies. A targeted factual consistency verification strategy is also established to ensure accurate reasoning.

• We evaluate LAMD on a collected dataset<sup>1</sup> that reflects the real-world setting. Our results show that LAMD outperforms conventional learning-based detectors under distribution shift, demonstrating its effectiveness in detecting and explaining fast evolving Android malware. Furthermore, we discuss the tension between learning-based models and pre-trained LLMs, highlighting promising directions for future research.

## 2. Related Work

This section reviews conventional Android malware detection, their real-world limitations, and recent advances in LLM-powered Android malware detection, situating our work in this evolving field.

## 2.1. Learning-based Android Malware Detection

Learning-based Android malware detectors leverage machine learning or deep learning models to automatically learn patterns from features extracted by static or dynamic analysis, enhancing scalability and adaptability [10]. Due to the high cost of dynamic analysis, most models rely on static feature extraction through reverse engineering [11], [12]. Despite advancements, these models struggle with concept drift, where evolving malware variants degrade performance in real-world deployments [1], [2]. Current works try to mitigate it, some of them focus on exploring robust features [13], [14], [15], [16] and others leverage continual learning [3], [17], [18] or active learning [19], [20] to let models adapt to new distribution. However, these methods either target to specific feature space [13] or introduce high retraining overhead and the risk of label poisoning [21], [22]. An additional challenge is explainability, which is critical for security analysis. Existing methods primarily use feature attribution techniques, providing importance scores without generating human-readable behavioral analysis [23], [24]. To address these limitations, researchers are increasingly exploring LLM-based approaches, which leverage extensive external knowledge and reasoning capabilities [25], [26], while tasks related to malware detection and analysis still remain under-explored.

### 2.2. LLM-powered Malware Detection

LLMs are increasingly used in security tasks like code analysis [27], vulnerability detection [28], [29], and malware classification [30]. Unlike traditional learning-based models, LLMs offer zero-shot inference capabilities, enabling them to generalize beyond predefined training data and feature spaces [4], [5]. This adaptability leads researchers to explore LLMs for malware detection and analysis, demonstrating

<sup>1.</sup> The dataset is open-source for further research on LLM-based malware tasks: https://doi.org/10.5281/zenodo.14884736



Figure 2: The workflow of LAMD. Suspicious APIs are identified via predefined rules (Step 1), and their calling functions with control flow graphs are extracted through static analysis. A customized backward slicing technique refines relevant instructions, preserving potential malicious intent (Step 2). In the code reasoning phase, the structured control flow graph, function relationships, and suspicious APIs form hierarchical tiers for malware detection and human-readable explanations (Steps 3-6). Factual consistency verification ensures first-tier summary reliability, mitigating hallucination (Step 4).

their potential in security tasks. However, most studies focus on relatively simple malware ecosystems, such as npm packages [31], PowerShell scripts [32], Linux binaries [30], and JavaScript-based threats [33].

Initial attempts at LLM-based Android malware detection are limited. Walton et al. [34] proposed a hierarchical approach, analyzing decompiled code at the function, class, and package levels. However, the lack of filtering mechanisms allows benign code to obscure malicious patterns, reducing detection accuracy and increasing computational costs. Even with a balanced 200-sample dataset, their best prompt only achieved 75% accuracy. Zhao et al. [35] relies on predefined feature spaces (Drebin [11]) to generate feature summaries instead of analyzing raw code. These summaries are then embedded and fed into a deep neural network (DNN) for training and detection. While focusing on the feature's name is efficient, it lacks structural and invocation insights. Additionally, as a learning-based method, it also inherits generalization issues and dataset bias influences of conventional models. These limitations highlight the need for a framework that leverages LLMs effectively in realworld Android malware detection, which integrates both structural and semantic context.

### 3. Methodology

This section outlines the core components of our framework, LAMD, and how they cooperate to detect and understand Android malware efficiently.

#### 3.1. Overall Architecture

The LAMD framework is designed to extract essential functionalities and their contextual information, enabling LLMs to generate both detection and reasoning results. It consists of two key components: (1) Key Context Extraction and (2) Tier-wise Code Reasoning, detailed as follows:

• **Key Context Extraction**: This module identifies suspicious APIs as seed points and analyzes their control and data dependencies within the application.

It provides a structured representation of key program behaviors by pruning the calling relationships of potentially malicious interactions.

• **Tier-wise Code Reasoning**: To preserve contextual integrity while managing token limitations, LAMD employs a tiered reasoning strategy. It processes information at three levels—function, API, and APK—where the output of each tier informs the next. To mitigate error propagation, factual consistency verification is applied at the first tier.

Figure 2 shows the pipeline of our framework. Overall, the raw input to LAMD consists of APK files, from which suspicious APIs and their sliced contexts are extracted as the input for LLMs. The malicious behavior of the application is then determined through three tiers of code reasoning.

### 3.2. Key Context Extraction

**3.2.1.** Suspicious API Collection. Malware exploits system vulnerabilities or API permissions to steal data, manipulate resources, or maintain persistence. Many attacks rely on sensitive API calls to implement malicious behaviors. We perform static analysis on APKs to extract suspicious APIs as key context to identify malware. Let  $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$  be the set of all API calls in an APK. A subset  $\mathcal{A}_{sus} \subset \mathcal{A}$  is deemed suspicious if it interacts with sensitive components, executes malware-associated operations or exposes sensitive data. These APIs fall into two categories:

• Sensitive data access APIs: Many apps handle sensitive data, but assessing developer trustworthiness is challenging. Monitoring APIs accessing such data is crucial. Smartphone OSs enforce permission-based access control, requiring declared permissions for some APIs, while others, like getPrimaryClip(), bypass enforcement. Therefore, an API  $a_i$  falls into this category if it either: (1) requires explicit permissions for access control, or (2) grants direct access to sensitive user data without permission enforcement. • Sensitive data transmission APIs: Monitoring potential data exfiltration channels is critical, as malware often exploits these APIs—commonly referred to as sink APIs—to transmit sensitive information to external entities. An API  $a_j$  is classified as suspicious if it facilitates the transfer of sensitive data to an external environment.

To extract suspicious APIs, we leverage publicly available knowledge based on PScout [36], SuSi [37] and Flowdroid [38] to label them. In a real-world setting, not every application contains suspicious APIs. These samples should be treated individually, where traditional learningbased malware detectors also struggle to handle [39].

**3.2.2. Backward Program Slicing.** Extracting suspicious APIs is useful, but analyzing them in isolation often obscures their context. For instance, while sendTextMessage() is legitimate in messaging apps, malware may exploit it for premium-rate SMS. To capture intent, we extract functions invoking suspicious APIs and refine their control flow graphs (CFGs),  $G = \langle N, E \rangle$ , where nodes  $n \in N$  represent basic blocks or instructions, and edges  $(e_1, e_2) \in E$  define control flow.

Recent work reveals that decoder-only Transformers, at the core of modern LLMs like GPT-4o-mini, are susceptible to over-squashing [40], where information from distant tokens is compressed and loses influence. The authors show how this may impair models in learning proper representations of similar long sequences, leading to representation clashes and wrong outcomes for downstream tasks. The study suggests that input-compressing strategy might preserve a high the signal-to-noise ratio, able to improve performance on downstream tasks. We follow this intuition in our context (the analysis of large complex programs), aiming at providing quality context to the models to improve on the downstream tasks. Since CFGs can be large and noisy, containing instructions that are irrelevant to suspicious API usages, directly inputting them to LLM risks conflating malicious and benign patterns, thereby obscuring sensitive behaviors. We therefore apply backward slicing [41] to isolate instructions affecting the API invocation and provide meaningful context to LLM. A slice S is defined by a slicing criterion  $C = \langle s, V \rangle$  where s is the statement invoking  $a_i$ and V includes all parameters. We classify relevant variables as: (1) Direct relevant variables: Variables' values can affect variable  $v \in V$  of  $a_i$  (2) Indirect relevant variables: Variables in branch statements whose value affects invocation of  $a_i$ . The backward slicing is to select the set of instructions in  $\mathcal{P}$  that directly or indirectly affect the execution or parameters of  $a_i$ . The backward slicing algorithm consists of two steps to ensure completeness in complex branch structures:

- Variable retrieval: Identify all variables contributing to the parameters (and internal states) used by the suspicious API and store them in a candidate set.
- Slices extraction: Append instructions related to variables collected in the first step.

After slicing, we generate sliced CFGs for each sensitive API, preserving essential control flow and statements. Notably, If undeclared variables remain in a sliced function, inter-procedural backward slicing is recursively applied to its callers until all variables are resolved. The details of slicing algorithm are shown in Appendix D.

#### 3.3. Tier-wise Code Reasoning

Code reasoning involves analyzing and interpreting code to understand its behavior, identify potential threats, and generate meaningful explanations. We propose a three-tier reasoning strategy that refines APK behavior analysis from fine- to coarse-grained levels, enhancing both prediction accuracy and interpretability. This hierarchical approach improves malicious component identification, mitigates LLM token-length limitations, and captures structural and invocation semantics through separable reasoning, ensuring a more effective and scalable malware detection framework.

**3.3.1. Tier 1: Function Behavior Summarization.** In the previous stage, several functions invoking suspicious APIs are extracted and sliced to maintain the related context. Each sliced CFG of the function is fed to LLM to capture low-level code patterns and functionalities.

Tier 1: Function Behavior Summarization Prompt

You are a cybersecurity expert specializing in Android malware analysis. Analyze the provided control flow graph including instructions related to sensitive API calls in detail. Control Flow Graph: {*CFG\_content*}

**3.3.2. Tier 2: API Intent Awareness.** The context of a specific API is typically determined by a series of functions. Beyond function invocation relationships, it is crucial to examine how inter-function associations influence the API's intent. Due to diverse contexts, an API may appear in multiple Function Call Graphs (FCGs). For instance, getDeviceId() is benign when used solely for local logging but becomes malicious when invoked within sendImeiToServer(), where it exfiltrates the IMEI to a remote server. Therefore, at this mid-tier, all functions associated with a suspicious API are structured into multiple FCGs to analyze its overall intent. Each node in an FCG is represented by the generated function summary in tier 1.

#### Tier 2: API Intent Awareness Prompt

You are a cybersecurity expert specializing in Android malware analysis. Analyze the main functionality and behavior of the provided sensitive API based on the function call graphs and a summary of each function's behavior. API name:  $\{API name\}$ API type:  $\{access/transfer\}$ for *i*-th Function Call Graph: FCG:  $\{FCG\_content\}$  $\{function\_name\}\{function\_summary\}$  **3.3.3. Tier 3: APK Maliciousness Judgement.** After extracting intents from suspicious APIs in an APK, LLMs assess its maliciousness and justify their decision. and generate Indicators of Compromise (IoCs), summarizing sensitive data access, external transmissions, and anomalous behavior to enhance transparency and trust.

#### Tier 3: APK Maliciousness Judgement Prompt

You are a cybersecurity expert specializing in Android malware analysis. Determine whether the application is MALWARE or BENIGN, citing indicators of compromise, evidence, and malicious patterns if present. Give a final prediction and key findings of your analysis. for *i*-th API: API name: {*API name*}

API type: {access/transfer} API intent: {API summary}

**3.3.4. Factual Consistency Verification.** Generating behavior summaries with LLMs risks hallucinations [42], producing facts inconsistent with instructions. To prevent error accumulation, we verify function-level summaries before higher-tier reasoning, leveraging their limited dependencies and concise structure.

Building on factual consistency verification [43], [44], we design a structured template to capture data dependencies in sliced CFGs. To enhance inference, we prompt the LLM to select corresponding data relationships from the input function based on specific definitions. We define five dependencies: variable-to-API interactions (direct, transitive, conditional) and inter-variable relationships (parallel, derived). The former tracks how variables influence API execution via assignments, call chains, and control flow, while the latter captures joint computation and derivation. Loop dependencies are excluded, as Soot expands loops when analysing binary code. Appendix C details these dependencies. For consistency verification, we integrate an incontext learning-based prompt with function summarization, querying the LLM to extract dependencies in the format: < dependencies type >:< variable names >.

#### Factual Consistency Verification Prompt

The provided control flow graph represents a slice of the function, identifying variable relationships for each statement leading to the final invocation statement that invokes  $\{function\_name\}$ . The output should follow the template:  $\{template\}$ Control Flow Graph:  $\{CFG\_content\}$ There are FIVE types of relationships: 1. Direct: Variables used directly as function parameters. *Example: invoker1.method(r2) \rightarrow r1, r2* 2. Transitive: Variables whose values flow through assignments but are not directly used in the invocation. *Example: r2 = r3.getValue()*; *invoker1.method(r2) \rightarrow r3*...

To evaluate the reliability of generated summaries, we

TABLE 1: Overall detection performance (%) across three test sets with increasing distribution drift (Test 1 < Test 2 < Test 3).

Model	,	Test 1	L		Test 2	2	Test 3			
	F1	FPR	FNR	F1	FPR	FNR	F1	FPR	FNR	
Drebin [11]	81.33	0.40	24.21	73.60	0.99	31.14	61.59	4.36	37.00	
DeepDrebin [12]	71.92	0.62	34.12	69.22	0.85	36.13	66.11	0.95	38.58	
Malscan [45]	66.37	0.73	46.83	65.45	0.84	47.73	64.54	0.91	48.30	
LAMD-R	75.34	5.39	15.24	75.71	4.79	12.99	75.83	4.84	11.19	
LAMD-F	87.63	2.00	10.37	87.28	1.85	9.74	87.21	1.77	9.83	
LAMD	90.24	1.26	8.44	90.16	1.38	7.79	89.85	1.30	8.47	

propose a Data Relationship Coverage (DRC) metric:

$$DRC = \frac{\#\{\text{correctly completed dependencies}\}}{\#\{\text{all selected dependencies}\}}.$$
 (1)

A summary is considered factually consistent if the LLM accurately reconstructs variable dependencies, i.e.,  $DRC \ge \theta$ , where  $\theta$  is a reliability threshold. Otherwise, the summary is revised to mitigate inaccuracies.

#### 4. Evaluation

This section presents a comprehensive evaluation of LAMD, assessing its Android malware detection performance in real-world scenarios and the quality of its generated explanations through a series of experiments.

Our experiments are conducted on an RTX A6000 GPU. For LLM-based reasoning, we utilize GPT-4o-mini [46], selected for its high efficiency, cost-effectiveness, and strong reasoning capabilities. The Data Relation Coverage (DRC) threshold  $\theta$  is set to 0.95 to balance computational efficiency and accuracy.

#### 4.1. Dataset Construction

To ensure a realistic dataset, we adhere to the following principles [3]: (1) Maintain temporal order in training and testing; (2) Preserve the real-world malware-to-benign ratio; (3) Ensure diversity by including packed, obfuscated, and varied market samples.

Based on these principles, we select Android APKs from Androzoo [47] according to their discovery time, which is determined by their submission to VirusTotal [48] (as release timestamps can be unreliable<sup>2</sup>). The dataset includes a decade of benign and malicious samples, labeled based on VirusTotal analysis from Androzoo. Samples flagged by more than four vendors are classified as malicious. To track malware family evolution, we used Euphony [49] to extract family labels. The dataset spans from 2014 to 2023<sup>3</sup>, comprising 13,794 samples. The remaining samples are evenly split into three test sets (about 3,015 each) to represent incremental distribution drift. Table 6 shows details about the used dataset.

<sup>2.</sup> Due to modifiable or randomly generated release times, some timestamps (dex date) are inaccurate.

<sup>3.</sup> as APK labels generally stabilize after about one year in the wild [50], we choose samples before 2024

TABLE 2: Out-of-distribution (OOD) detection performance (%) on post-cutoff samples for LAMD and aligned training-temporal scope for learning-based models.

Model	F1	FPR	FNR
Drebin [11]	83.72	0.42	22.02
DeepDrebin [12]	79.78	0.39	26.80
Malscan [45]	83.87	0.00	46.77
LAMD	85.71	1.89	14.29

#### 4.2. Metrics

(1) Classification Metrics. To address class imbalance in Android malware datasets, we use the F1-score to balance precision and recall, while also minimizing False Positive Rate (FPR) and False Negative Rate (FNR) to improve accuracy and reduce manual analysis overhead. Results are reported as percentages. (2) Summarization Metrics. Effective malware analysis provides interpretable insights and aids manual audits. Since malware family identification often requires expert review, initial categorization prioritizes the common sense of behavior patterns. Following prior work [51], we mainly consider six categories: Adware, Backdoor, PUA (Potentially Unwanted Applications), Riskware, Scareware, and Trojan. For evaluation, we adopt a ChatGPT-based metric [52], [53], [54], where GPT-4omini [46] assesses whether LAMD's detection aligns with the expected behaviors of each category [55].

#### 4.3. Baselines

In Android malware detection, an APK serves as input, containing the codebase (e.g., dex files) and configuration files (e.g., AndroidManifest.xml), which provide behavioral insights like API calls and permissions, represented in vector or graph formats. We evaluate LAMD against Drebin [11], DeepDrebin [12], and Malscan [45] which are representative learning-based methods on these feature formats, with details in Appendix A. To assess component impact, LAMD-R removes tier-wise reasoning to test structural and semantic analysis, while LAMD-F retains reasoning but omits factual consistency verification to evaluate hallucination control.

#### 4.4. Experiments

We conduct a series of experiments to thoroughly evaluate LAMD's performance in malware detection and the quality of its generated explanation.

**4.4.1. Malware Detection Performance.** Table 1 compares LAMD's detection performance with baselines on test datasets with the increasing level of drift. LAMD improves F1-scores by 23.12% and reduces FNR by 71.59% on average, enhancing detection reliability. While the FPR shows a slight increase, it is less indicative of true performance due to class imbalance, where learning-based methods often misclassify malware as benign because of the dominance

TABLE 3: The quality evaluation of generated detection analysis. The "Correct" column indicates the number of categories correctly classified by the LLM.

Category	Family	Total	Correct
Adware	gexin, ewind	20	15
Backdoor	mobby, hiddad	20	18
PUA	umpay, scamapp, apptrack	13	9
Riskware	jiagu, smspay, smsreg	25	18
Scareware	fakeapp	10	9
Trojan	hqwar, hypay	12	12
Overall	/	100	81

of benign samples. The performance drop in LAMD-R highlights the necessity of hierarchical code summarization, and the slight decline in LAMD-F underscores the role of hallucination mitigation. By refining input code and extracting key information, LAMD minimizes hallucination risks, ensuring more reliable malware detection.

We performed an additional out-of-distribution (OOD) experiments designed to simulate a temporal natural drift by aligning GPT-4o-mini's pre-training cutoff date (October 2023) with a dataset comprising samples of newer timestamps. As detailed in Table 5, to make a fair comparison, LAMD is tested on samples from November and December 2023, while learning-based models are assessed from May-June 2020, the first two months following their training phase. Table 2 confirms the trend we have seen in Table 1: LAMD outperforms baselines and state-of-the-art methods even under this fair settings. We also observe notable tension between the two paradigms: the breakdown of F1 score reports a higher FPR and a lower FNR for LAMD compared to learning-based approaches. We hypothesize that zero-shot learners such as LAMD have better generalizability benefit from their large-scale pre-trained dataset that may generalize better against similar threats (resulting in lower FNR). Conversely, they suffer from making more mistakes due to the lack of a smaller, task-dependent dataset (lower FPR for learning-based approaches). If true, this opens interesting research opportunities to study tensions between the two paradigms.

**4.4.2. Effectiveness of explanations.** To assess analysis quality, we validate 100 correctly detected malware samples. Results show that 81 out of 100 samples are correctly classified into their respective categories. Table 3 summarizes the sample distribution and classification accuracy across categories. Due to their less distinct malicious patterns, Adware and Riskware pose greater challenges to accurate analysis compared to other categories, contributing to their higher misclassification rates.

**4.4.3. Cost.** We estimate the overall cost of evaluating LAMD. Using GPT-4o-mini as our base LLM, which costs \$0.15 per 1M tokens for input and \$0.6 per 1M tokens for output [46], the total expense for our test set of 9,046 APKs amounts to approximately \$1800, averaging \$0.199 per apk.

The cost for each apk varies depending on the number of sliced CFGs extracted and suspicious APIs the app invokes. While this demonstrates the feasibility of deploying LAMD in real-world scenarios, the operational expense remains non-trivial for large-scale applications. Future work could explore the effectiveness of applying LAMD with locally deployed LLMs like DeepSeek [56] and Llama [57] to avoid cloud-based API fees.

## 5. Case Study

This section demonstrates how LAMD overcomes the limitations of current LLMs in Android malware detection. Since LLMs cannot process APKs directly, we use JADX<sup>4</sup> to decompile them, concatenating all pseudo source codes for input [52]. Besides GPT-40-mini, we select Gemini 1.5 pro [6] as another comparison, which claims their longest context windows and for its malware detection capabilities.

Figure 1 highlights key failure cases. In the first example, we decompile a randomly selected sample<sup>5</sup>, which contains 1,547,806 lines of code, even after restricting the scope to the "source/com" folder containing critical files. Gemini 1.5 Pro fails to process it, exceeding its 20,971,520-byte context limit, while LAMD enables GPT-40-mini (with 10x lower token capacity) to generate an accurate prediction, demonstrating its ability to efficiently analyze large-scale Android applications within constrained LLM contexts.

In the second case, we verify the detection ability of LAMD. Analyzing an SMSReg malware sample<sup>6</sup>, Gemini misclassifies it as "BENIGN", erroneously identifying it as part of the "*xUtils*" library due to excessive benign classes disguising malicious behavior. With LAMD, both GPT-4o-mini and Gemini 1.5 pro correctly classify it as malware and identify key behaviors. The SMSReg malware family usually harvests device data, sends unauthorized SMS, and registers users for premium services to evading detection. Detection explanations enhanced by LAMD align with these behaviors, showing unauthorized access to getDeviceId(), getSubscriberId(), and sendTextMessage(), reflection-based evasion, insecure SSL handling, and excessive location tracking—all indicative of SMS fraud and unauthorized transactions.

# GPT-40-mini Output with LAMD:

Final Prediction: MALWARE Key Findings: \* Sensitive API Misuse: Uses "getDeviceId()", "getSubscriberId()", and "sendTextMessage()" via reflection to bypass access controls for data harvesting. \* SSL Vulnerability: Unchecked "SslErrorHandler.proceed()" and "onReceivedSslError" expose users to MITM attacks. \* Location Tracking: Collects "getLastKnownLocation" and "requestLocationUpdates" without user consent.

4. https://github.com/skylot/jadx

5. MD5: c37e223e3388b31b323ad39af45180fc

6. MD5: 2be97287c6af70f2074686b1a9021c06

## 6. Discussion

While LLMs improve Android malware detection, their general pre-training limits fine-grained behavior analysis. The phenomenon of over-squashing [40] in LLMs still introduces challenges to successfully identify key aspects within complex programs. Future research can benefit from domain-specific fine-tuning or external knowledge integration [58]. Meanwhile, learning-based methods remain essential for capturing overall insights in non-drift scenarios. Our experiments find that while LLMs excel in adapting to unseen threats, they require significantly more computational resources and incur higher costs, highlighting effective strategies to combine these two paradigms and maximize their complementary strengths. In future work, we plan to explore how LAMD and more generally context-aware LLMs perform against recent approaches aimed at learning temporally invariant features [59], [60]. These findings open interesting research opportunities to better understand the tension between learning-based models and pre-trained LLMs.

### 7. Conclusion

Large language models' superior zero-shot inference offers a promising solution for Android malware detectors to handle distribution drift, dataset bias and explainability gaps in real-world scenarios but struggle with excessive support code and complex program structures. To address these challenges, we propose LAMD, the first practical framework enabling LLMs for explainable Android malware detection. Our evaluation in the real-world setting demonstrates that LAMD outperforms conventional detectors by effectively analyzing complex structures and semantics. LAMD unlocks LLMs' potential in Android security, paving the way for AIdriven malware analysis.

### References

- R. Jordaney, K. Sharad, S. Dash, Z. Wang, D. Papini, I. Nouretdinov, and L. Cavallaro, "Transcend: Detecting concept drift in malware classification models," USENIX Security Symposium, USENIX Security Symposium, Dec 2016.
- [2] F. Barbero, F. Pendlebury, F. Pierazzi, and L. Cavallaro, "Transcending transcend: Revisiting malware classification in the presence of concept drift," in 2022 IEEE Symposium on Security and Privacy (SP). IEEE, 2022, pp. 805–823.
- [3] F. Pendlebury, F. Pierazzi, R. Jordaney, J. Kinder, and L. Cavallaro, "{TESSERACT}: Eliminating experimental bias in malware classification across space and time," in 28th USENIX security symposium (USENIX Security 19), 2019, pp. 729–746.
- [4] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," Advances in neural information processing systems, vol. 35, pp. 22 199–22 213, 2022.
- [5] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom, "Toolformer: Language models can teach themselves to use tools," *Advances in Neural Information Processing Systems*, vol. 36, pp. 68 539–68 551, 2023.
- [6] G. Team, P. Georgiev, V. I. Lei, R. Burnell, L. Bai, A. Gulati, G. Tanzer, D. Vincent, Z. Pan, S. Wang *et al.*, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," *arXiv* preprint arXiv:2403.05530, 2024.
- [7] X. Zhou, K. Kim, B. Xu, D. Han, and D. Lo, "Out of sight, out of mind: Better automatic vulnerability repair by broadening input ranges and sources," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.
- [8] A. Hindle, E. T. Barr, M. Gabel, Z. Su, and P. Devanbu, "On the naturalness of software," *Communications of the ACM*, vol. 59, no. 5, pp. 122–131, 2016.
- [9] H. He, X. Lin, Z. Weng, R. Zhao, S. Gan, L. Chen, Y. Ji, J. Wang, and Z. Xue, "Code is not natural language: Unlock the power of semantics-oriented graph representation for binary code similarity detection," in *33rd USENIX Security Symposium (USENIX Security* 24), PHILADELPHIA, PA, 2024.
- [10] J. Liu, J. Zeng, F. Pierazzi, L. Cavallaro, and Z. Liang, "Unraveling the key of machine learning solutions for android malware detection," *arXiv preprint arXiv*:2402.02953, 2024.
- [11] D. Arp, M. Spreitzenbarth, M. Hubner, H. Gascon, K. Rieck, and C. Siemens, "Drebin: Effective and explainable detection of android malware in your pocket," in *Ndss*, vol. 14, 2014, pp. 23–26.
- [12] K. Grosse, N. Papernot, P. Manoharan, M. Backes, and P. McDaniel, "Adversarial examples for malware detection," in *Computer Security–ESORICS 2017: 22nd European Symposium on Research in Computer Security, Oslo, Norway, September 11-15, 2017, Proceedings, Part II 22.* Springer, 2017, pp. 62–79.
- [13] Y. Yang, B. Yuan, J. Lou, and Z. Qin, "Scrr: Stable malware detection under unknown deployment environment shift by decoupled spurious correlations filtering," *IEEE Transactions on Dependable and Secure Computing*, 2024.
- [14] N. McLaughlin, J. Martinez del Rincon, B. Kang, S. Yerima, P. Miller, S. Sezer, Y. Safaei, E. Trickel, Z. Zhao, A. Doupé *et al.*, "Deep android malware detection," in *Proceedings of the seventh ACM on conference on data and application security and privacy*, 2017, pp. 301–308.
- [15] D. Angioni, L. Demetrio, M. Pintor, and B. Biggio, "Robust machine learning for malware detection over time," *arXiv preprint arXiv:2208.04838*, 2022.
- [16] X. Zhang, Y. Zhang, M. Zhong, D. Ding, Y. Cao, Y. Zhang, M. Zhang, and M. Yang, "Enhancing state-of-the-art classifiers with api semantics to detect evolved android malware," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer* and Communications Security, Oct 2020. [Online]. Available: http://dx.doi.org/10.1145/3372297.3417291

- [17] Y. Chen, Z. Ding, and D. Wagner, "Continuous learning for android malware detection," in 32nd USENIX Security Symposium (USENIX Security 23), 2023, pp. 1127–1144.
- [18] V. Y. F. Tan, P. L.A., and K. Jagannathan, "Recent advances in concept drift adaptation methods for deep learning," in *Proceedings of the Thirty-First International Joint Conference* on Artificial Intelligence, Jun 2022. [Online]. Available: http://dx.doi.org/10.24963/ijcai.2022/784
- [19] K. Xu, Y. Li, R. Deng, K. Chen, and J. Xu, "Droidevolver: Selfevolving android malware detection system," in 2019 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE, 2019, pp. 47–62.
- [20] F. Ceschin, M. Botacin, H. M. Gomes, F. Pinagé, L. S. Oliveira, and A. Grégio, "Fast & furious: On the modelling of malware detection as an evolving data stream," *Expert Systems with Applications*, vol. 212, p. 118590, 2023.
- [21] S. Yang, X. Zheng, J. Li, J. Xu, X. Wang, and E. C. Ngai, "Recda: Concept drift adaptation with representation enhancement for network intrusion detection," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 3818–3828.
- [22] Z. Kan, F. Pendlebury, F. Pierazzi, and L. Cavallaro, "Investigating labelless drift adaptation for malware detection," in *Proceedings of the 14th ACM Workshop on Artificial Intelligence and Security*, 2021, pp. 123–134.
- [23] Y. He, J. Lou, Z. Qin, and K. Ren, "Finer: Enhancing state-of-the-art classifiers with feature attribution to facilitate security analysis," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 416–430.
- [24] M. K. Belaid, E. Hüllermeier, M. Rabus, and R. Krestel, "Do we need another explainable ai method? toward unifying post-hoc xai evaluation methods into an interactive and multi-dimensional benchmark," *arXiv preprint arXiv*:2207.14160, 2022.
- [25] S. Sarsa, P. Denny, A. Hellas, and J. Leinonen, "Automatic generation of programming exercises and code explanations using large language models," in *Proceedings of the 2022 ACM Conference on International Computing Education Research-Volume 1*, 2022, pp. 27–43.
- [26] S. MacNeil, A. Tran, D. Mogil, S. Bernstein, E. Ross, and Z. Huang, "Generating diverse code explanations using the gpt-3 large language model," in *Proceedings of the 2022 ACM Conference on International Computing Education Research-Volume 2*, 2022, pp. 37–39.
- [27] T. Ahmed, K. S. Pai, P. Devanbu, and E. Barr, "Automatic semantic augmentation of language model prompts (for code summarization)," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.
- [28] Y. Sun, D. Wu, Y. Xue, H. Liu, W. Ma, L. Zhang, Y. Liu, and Y. Li, "Llm4vuln: A unified evaluation framework for decoupling and enhancing llms' vulnerability reasoning," *arXiv preprint* arXiv:2401.16185, 2024.
- [29] S. Ullah, M. Han, S. Pujar, H. Pearce, A. Coskun, and G. Stringhini, "Llms cannot reliably identify and reason about security vulnerabilities (yet?): A comprehensive evaluation, framework, and benchmarks," in *IEEE Symposium on Security and Privacy*, 2024.
- [30] P. M. S. Sánchez, A. H. Celdrán, G. Bovet, and G. M. Pérez, "Transfer learning in pre-trained large language models for malware detection based on system calls," arXiv preprint arXiv:2405.09318, 2024.
- [31] Z. Yu, M. Wen, X. Guo, and H. Jin, "Maltracker: A fine-grained npm malware tracker copiloted by llm-enhanced dataset," in *Proceedings* of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis, 2024, pp. 1759–1771.
- [32] J. Deng, X. Li, Y. Chen, Y. Bai, H. Weng, Y. Liu, T. Wei, and W. Xu, "Raconteur: A knowledgeable, insightful, and portable llm-powered shell command explainer," *arXiv preprint arXiv:2409.02074*, 2024.

- [33] Y. Zhang, X. Zhou, H. Wen, W. Niu, J. Liu, H. Wang, and Q. Li, "Tactics, techniques, and procedures (ttps) in interpreted malware: A zero-shot generation with large language models," *arXiv preprint* arXiv:2407.08532, 2024.
- [34] B. J. Walton, M. E. Khatun, J. M. Ghawaly, and A. Ali-Gombe, "Exploring large language models for semantic analysis and categorization of android malware," *arXiv preprint arXiv:2501.04848*, 2025.
- [35] W. Zhao, J. Wu, and Z. Meng, "Apppoet: Large language model based android malware detection via multi-view prompt engineering," *Expert Systems with Applications*, vol. 262, p. 125546, 2025.
- [36] K. W. Y. Au, Y. F. Zhou, Z. Huang, and D. Lie, "Pscout: analyzing the android permission specification," in *Proceedings of the 2012* ACM conference on Computer and communications security, 2012, pp. 217–228.
- [37] S. Rasthofer, S. Arzt, and E. Bodden, "A machine-learning approach for classifying and categorizing android sources and sinks." in *NDSS*, vol. 14, no. 1125, 2014.
- [38] S. Arzt, S. Rasthofer, C. Fritz, E. Bodden, A. Bartel, J. Klein, Y. Le Traon, D. Octeau, and P. McDaniel, "Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps," ACM sigplan notices, vol. 49, no. 6, pp. 259–269, 2014.
- [39] Y. He, Y. Liu, L. Wu, Z. Yang, K. Ren, and Z. Qin, "Msdroid: Identifying malicious snippets for android malware detection," *IEEE Transactions on Dependable and Secure Computing*, vol. 20, no. 3, pp. 2025–2039, 2022.
- [40] F. Barbero, A. Banino, S. Kapturowski, D. Kumaran, J. Madeira Araújo, O. Vitvitskyi, R. Pascanu, and P. Veličković, "Transformers need glasses! information over-squashing in language tasks," Advances in Neural Information Processing Systems, vol. 37, pp. 98 111–98 142, 2024.
- [41] M. Weiser, "Program slicing," IEEE Transactions on software engineering, no. 4, pp. 352–357, 1984.
- [42] S. Farquhar, J. Kossen, L. Kuhn, and Y. Gal, "Detecting hallucinations in large language models using semantic entropy," *Nature*, vol. 630, no. 8017, pp. 625–630, 2024.
- [43] Y. Li, L. Li, Q. Yang, M. Litvak, N. Vanetik, D. Hu, Y. Li, Y. Zhou, D. Xu, and X. Zhang, "Just cloze! a fast and simple method for evaluating the factual consistency in abstractive summarization," arXiv preprint arXiv:2210.02804, 2022.
- [44] A. Wang, K. Cho, and M. Lewis, "Asking and answering questions to evaluate the factual consistency of summaries," *arXiv preprint* arXiv:2004.04228, 2020.
- [45] Y. Wu, X. Li, D. Zou, W. Yang, X. Zhang, and H. Jin, "Malscan: Fast market-wide mobile malware scanning by social-network centrality analysis," in 2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE). IEEE, 2019, pp. 139– 150.
- [46] OpenAI, "Hello gpt-4o," 2024, accessed: Feb. 3, 2025. [Online]. Available: https://openai.com/index/gpt-4o-mini-advancingcost-efficient-intelligence/
- [47] K. Allix, T. F. Bissyandé, J. Klein, and Y. Le Traon, "Androzoo," in *Proceedings of the 13th International Conference* on Mining Software Repositories, May 2016. [Online]. Available: http://dx.doi.org/10.1145/2901739.2903508
- [48] VirusTotal, "Virustotal: Free online virus, malware and url scanner," 2004, accessed: 2025-02-14. [Online]. Available: https://www.virustotal.com
- [49] M. Hurier, G. Suarez-Tangil, S. K. Dash, T. F. Bissyandé, Y. L. Traon, J. Klein, and L. Cavallaro, "Euphony: harmonious unification of cacophonous anti-virus vendor labels for android malware," in *Proceedings of the 14th International Conference on Mining Software Repositories*. IEEE Press, 2017, pp. 425–435.

- [50] B. Miller, A. Kantchelian, M. C. Tschantz, S. Afroz, R. Bachwani, R. Faizullabhoy, L. Huang, V. Shankar, T. Wu, G. Yiu *et al.*, "Reviewer integration and performance measurement for malware detection," in *Detection of Intrusions and Malware, and Vulnerability Assessment: 13th International Conference, DIMVA 2016, San Sebastián, Spain, July 7-8, 2016, Proceedings 13.* Springer, 2016, pp. 122–141.
- [51] I. W. Canada, "Understanding android malware families (uamf): The foundations (article 1)," 2024, accessed: February 6, 2025. [Online]. Available: https://www.itworldcanada.com/blog/understandingandroid-malware-families-uamf-the-foundations-article-1/441562
- [52] C. Fang, N. Miao, S. Srivastav, J. Liu, R. Zhang, R. Fang, R. Tsang, N. Nazari, H. Wang, H. Homayoun *et al.*, "Large language models for code analysis: Do {LLMs} really do their job?" in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 829–846.
- [53] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg *et al.*, "Sparks of artificial general intelligence: Early experiments with gpt-4. arxiv," *arXiv preprint arXiv:2303.12712*, 2023.
- [54] C.-H. Chiang and H.-y. Lee, "Can large language models be an alternative to human evaluations?" arXiv preprint arXiv:2305.01937, 2023.
- [55] M. Gao, J. Ruan, R. Sun, X. Yin, S. Yang, and X. Wan, "Human-like summarization evaluation with chatgpt," *arXiv preprint* arXiv:2304.02554, 2023.
- [56] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi *et al.*, "Deepseek-r1: Incentivizing reasoning capability in Ilms via reinforcement learning," *arXiv preprint arXiv:2501.12948*, 2025.
- [57] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan *et al.*, "The llama 3 herd of models," *arXiv preprint arXiv:2407.21783*, 2024.
- [58] Y. He, H. She, X. Qian, X. Zheng, Z. Chen, Z. Qin, and L. Cavallaro, "On benchmarking code llms for android malware analysis," *arXiv* preprint arXiv:2504.00694, 2025.
- [59] X. Zheng, S. Yang, E. C. Ngai, S. Jana, and L. Cavallaro, "Learning temporal invariance in android malware detectors," *arXiv preprint* arXiv:2502.05098, 2025.
- [60] Y. He, J. Lei, Z. Qin, and K. Ren, "Dream: Combating concept drift with explanatory detection and adaptation in malware classification," arXiv preprint arXiv:2405.04095, 2024.
- [61] M. L. Menéndez, J. Pardo, L. Pardo, and M. Pardo, "The jensenshannon divergence," *Journal of the Franklin Institute*, vol. 334, no. 2, pp. 307–318, 1997.

# Appendix A. Baseline Models

Drebin [11] detects malware using binary feature vectors derived from nine data types (e.g., hardware components, API calls, permissions) and classifies samples via a linear classifier.

DeepDrebin [12] extends Drebin by replacing the linear classifier with a three-layer deep neural network (DNN) while retaining the same feature space for feature extraction and classification.

Malscan [45] employs a graph-based approach, extracting sensitive API calls from APKs and computing four centrality measures (degree, Katz, proximity, and harmonic wave centralities) as features. We use the optimal feature with a Random Forest classifier, which leads to the lowest overhead.

All of them can be considered as the current state-ofthe-art in Android malware detection.

TABLE 4: Overview of the evaluation dataset, where M denotes malware and B denotes benign applications.

Test Set	Time Interval	Sample Size	<b>Existing Families</b>	New Families	Packed	Malicious (M)	Benign (B)	M/(M+B) %
Test Set 1	2020.05 - 2021.01	3015	21	24	18	284	2731	9.42
Test Set 2	2021.01 - 2021.12	3015	28	32	30	298	2717	9.88
Test Set 3	2021.12 - 2023.12	3016	34	36	40	302	2714	10.01

TABLE 5: Overview of the OOD evaluation dataset aligned with the GPT-40-mini pre-training cutoff (October 2023), where M denotes malware and B denotes benign applications.

Test Set	Time Interval	Sample Size	<b>Existing Families</b>	New Families	Packed	Malicious (	M) Benign (B	) M/(M+B) %
Learning-based	2020.05 - 2020.06	884	15	7	2	62	822	7.01
LAMD	2023.11 - 2023.12	60	0	1	0	7	53	11.67

TABLE 6: Performance of learning-based methods using Drebin features under different training/validation data volumes.

Sample Size	Model	V	Validation			Test 1			Test 2			Test 3		
~~~~		F1	FPR	FNR	F1	FPR	FNR	F1	FPR	FNR	F1	FPR	FNR	
13k	Drebin	94.02	0.57	6.58	81.33	0.40	24.21	73.60	0.99	31.14	61.59	4.36	37.00	
	DeepDrebin	95.72	0.24	6.06	71.92	0.62	34.12	69.22	0.85	36.13	66.11	0.95	38.58	
60k	Drebin	95.83	0.34	5.28	68.76	0.68	36.67	63.41	1.91	40.12	62.83	2.33	40.28	
	DeepDrebin	96.68	0.28	4.15	68.82	0.64	36.65	63.48	1.72	40.17	58.99	4.31	41.19	

## Appendix B. Impact of Training Dataset Sizes

In real-world scenarios, a large number of Android applications are packed and obfuscated to bypass malware detection systems and complicate manual analysis. A notable shift occurred post-2020, as more Androzoo samples adopted these techniques. Consequently, our test set, drawn from 2020, reflects a more realistic setting. Notably, the Jensen–Shannon (JS) divergence [61] (measuring the similarity between two probability distributions) between the training and test sets is five times larger than that between the training and validation sets.

Intuitively, increased training samples should enhance detection performance in learning-based methods. While LAMD excels in zero-shot learning, one may argue that traditional methods are hindered by smaller training datasets compared to the vast pretraining data leveraged by LLMs. However, the impact of significant drift warrants reconsideration. Table 6 highlights the performance of learning-based detectors (e.g., Drebin feature space) across varying dataset sizes, across the training dataset timeframe. The results indicate that additional training data does not improve detection of samples with significant drift, though it enhances performance on in-distribution samples.

Expanding the training set introduces more information, helping detectors learn clearer decision boundaries, consistent with the principle of empirical risk minimization (ERM). However, ERM identifies features that distinguish samples across the entire dataset, and when the training set is heavily skewed toward older data (e.g., 2014-2019/3), the model is more likely to learn features representative of early apps rather than more recent ones. Consequently, when tested on subsequent months (e.g., Test 1), the model exhibits worse performance, as it relies on features that are less effective for newer, drifted samples. This suggests that while increasing training data is generally beneficial, it may reinforce outdated patterns in the presence of significant drift, ultimately impairing generalizability.

# Appendix C. Factual Consistency Verification

To check factual consistency, we leverage data dependencies that capture relationships between variables and APIs. We focus on five dependencies (see Table 7) as they represent fundamental program relationships critical for malware reasoning and detection. Specifically:

(1) Variable-to-API Dependencies. These dependencies determine whether the execution of an API is influenced by specific variables.

- Direct dependencies: ensure that variables explicitly control API calls, providing strong evidence of intended execution.
- Transitive dependencies: track how variables propagate through function calls.
- Conditional dependencies: account for control-flow influences (e.g., if statements), identifying cases where malicious logic may be context-dependent.

(2) Inter-variable Dependencies. These capture relationships between variables that may affect security-sensitive operations.

• Parallel dependencies: detect multiple variables jointly contributing to a computation, highlighting complex conditions leading to API execution.

TABLE 7: Data dependencies used for factual consistency verification.

Category	Туре	Description
Variable-to-API	Direct Transitive Conditional	A variable directly determines API execution. A variable is propagated through a call chain to an API. A variable influences API execution via control flow (e.g., if/else).
Inter-variable	Parallel Derived	Two variables jointly contribute to computing another variable. A variable is derived from another through computation.

• Derived dependencies: reveal computations that transform one variable into another, helping detect disguised or obfuscated malicious behaviors.

By incorporating these dependencies into factual consistency verification, LAMD ensures that summarized malicious behaviors are not based on hallucinations but are grounded in actual program logic, improving both precision and interpretability in Android malware detection.

## Appendix D. Backward Slicing Algorithm

The proposed backward slicing algorithm is shown in Algorithm 1. This backward slicing procedure consists of two phases: variable retrieval and slice extraction. In Stage 1, the algorithm tracks variables that influence the suspicious API invocation  $a_i$ . Specifically, in Line 2, the algorithm initializes the *worklist* with *unit* containing  $a_i$ , and in Line 4, it begins iterating through the control flow graph (CFG) in a backward manner starting from *unit*, extracting relevant variables from each unit. The algorithm updates the locals list by removing defined variables and adding those that affect  $a_i$  in Line 13, and stores the updated variables in varMap in Line 15. Stage 2 focuses on extracting slices based on varMap. In Line 23, the algorithm resets the worklist and visited sets, then starts from the predecessors of the original unit that invoked  $a_i$ . Line 28 checks the successors of each unit and identifies those relevant to the API invocation. The combined slices are formed in Line 29, integrating all the relevant variables and units, and the algorithm returns the slices in Line 38. This two-phase process enables the algorithm to isolate the parts of the function that influence the suspicious API invocation, facilitating further analysis.

#### Algorithm 1 Algorithm of Backward Slicing

**Require:** Control Flow Graph CFG of the function that invokes the suspicious API  $a_i$  at unit block *unit* 

- 1: Stage 1: Variable Retrieval
- 2: Initialize *worklist* with *unit* containing  $a_i$
- 3: Initialize  $visited \leftarrow \{\}, varMap \leftarrow \emptyset$
- 4: while  $worklist \neq \emptyset$  do
- 5: Pop *currUnit* from *worklist*
- 6: Initialize  $locals \leftarrow \{\}$
- 7: **if** currUnit contains  $a_i$  **then**
- 8: Extract variables used by  $a_i$  and add to *locals*
- 9: **else**
- 10: Load variables from successors of *currUnit* using *varMap* and add to *locals*
- 11: end if
- 12: **if** currUnit is relevant to the  $a_i$  invocation based on *locals* **then**
- 13: Update *locals*: remove defined variables, add used variables that affect the  $a_i$  invocation
- 14: **end if**
- 15: Store *locals* in *varMap*[*currUnit*]
- 16: for  $pred \in$  predecessors of currUnit do
- 17: **if**  $pred \notin visited$  **then**
- 18: Add predecessor *pred* to *visited* and *worklist*
- 19: **end if**
- 20: end for
- 21: end while
- 22: Stage 2: Slices Extraction
- 23: Initialize  $slices \leftarrow \{unit\}$ , reset worklist and visited
- 24: Add predecessors of unit to worklist
- 25: while  $worklist \neq \emptyset$  do
- 26: Pop *currUnit* from *worklist*
- 27: Load variables from the successors of *currUnit* using *varMap*
- 28: **if** currUnit is relevant to the  $a_i$  invocation based on the successor variable **then**
- 29: Add *currUnit* to *slices*
- 30: end if
- 31: for  $pred \in$  predecessors of currUnit do
- 32: **if**  $pred \notin visited$  **then**
- 33: Add predecessor *pred* to *visited* and *worklist*
- 34: **end if**
- 35: **end for**
- 36: end while
- 37:
- 38: return *slices*