Rethinking and Exploring String-Based Malware Family Classification in the Era of LLMs and RAG

Yufan Chen[1], Daoyuan Wu[2*], Juantao Zhong[2], Zicheng Zhang[3], Debin Gao[3], Shuai Wang[2], Yingjiu Li[4], and Ning Liu[1,5]

[1] CityU Shenzhen Research Institute, yfchen2224@cityu.edu.cn

[2] The Hong Kong University of Science and Technology, {daoyuan|juantao|shuaiw}@cse.ust.hk

[3] Singapore Management University, zczhang.2020@phdcs.smu.edu.sg, dbgao@smu.edu.sg

[4] University of Oregon, yingjiul@uoregon.edu

[5] City University of Hong Kong, ninliu@cityu.edu.hk

***Abstract*—Malware Family Classification (MFC) aims to identify the fine-grained family (e.g., `GuLoader` or `BitRAT`) to which a potential malware sample belongs, in contrast to malware detection or sample classification that predicts only an Yes/No. Accurate family identification can greatly facilitate automated sample labeling and understanding on crowdsourced malware analysis platforms such as VirusTotal and MalwareBazaar, which generate vast amounts of data daily. In this paper, we explore and assess the feasibility of using traditional binary string features for MFC in the new era of large language models (LLMs) and Retrieval-Augmented Generation (RAG). Specifically, we investigate how *Family-Specific String* (FSS) features could be utilized in a manner similar to RAG to facilitate MFC. To this end, we develop a curated evaluation framework covering 4,347 samples from 67 malware families, extract and analyze over 25 million strings, and conduct detailed ablation studies to assess the impact of different design choices in four major modules.**

I. Introduction

Malware family classification (MFC) plays a pivotal role in threat intelligence [35], [42] and malware analysis pipelines [5], [6]. Rather than merely determining whether a sample is malicious, MFC aims to assign a specific family label—such as `GuLoader` and `BitRAT`—that reflects common behavioral traits, infection vectors, and capabilities. This fine-grained family attribution is essential for tracking malware evolution, generating automated YARA rules, and supporting downstream tasks like clustering, triage, and prioritization. In particular, crowdsourced malware analysis platforms such as VirusTotal [6] and MalwareBazaar [5] require accurate family identification to facilitate automated sample labeling and understanding on their vast amounts of data daily.

Prior research has explored various strategies for malware classification, ranging from manual signature engineering [36] and behavior-based heuristics [16], [31] to learning representations from structural graphs [44], disassembled code [34], and API call traces [16], [31], [46]. However, despite their potential, these approaches often suffer from significant limitations: handcrafted features do not generalize well [8], dynamic analysis [48], [50] is expensive and can be evaded [19], and learned representations—particularly those derived from binary content—tend to be opaque, hard to interpret, and vulnerable to adversarial obfuscation [49]. As a result, most recent studies treat strings merely as auxiliary inputs; see §II.

In this paper, we rethink a long-standing but underutilized resource in malware analysis: string artifacts extracted from binaries. These strings can contain human-readable indicators such as command-line options, configuration paths, encryption routines, URLs, domain names, registry keys, and API call names—features that, when properly filtered and semantically understood, can offer rich insights into malware behavior. Historically, string features have been overshadowed by more complex modalities such as control flow or dynamic behaviors. However, we argue that the rise of large language models (LLMs) [30], [40], particularly when combined with retrieval-augmented generation (RAG) [22], presents a new opportunity to revisit string-based malware family classification.

Our work proposes an exploratory pipeline centered on a new feature abstraction: *Family-Specific Strings (FSS)*. We define FSS as strings that appear within a specific malware family but are absent in other families. By curating a database of these strings from a labeled training set, embedding them into a semantic vector space, and matching them against features extracted from new samples, we aim to *explore and assess the feasibility of using traditional binary string features for MFC in the new era of LLMs and RAG*. To systematically examine this vision, we design a four-stage pipeline and conduct an extensive empirical study evaluating each component in isolation and combination. Our goal is not to propose a fixed system but rather to *explore the design space of FSS-centered MFC through four key research questions*.

The first question (RQ1) concerns string extraction: can static-only extraction methods, such as FLOSS [4], recover enough meaningful strings for robust classification, or do we need dynamic execution (e.g., sandbox-based memory snapshots) to supplement the feature space? Our analysis reveals that static extraction alone leaves many semantically rich strings hidden—particularly for heavily obfuscated families like `Stop` or `Formbook`—and that hybrid approaches can improve classification accuracy by up to 400% in some cases. This underscores the importance of runtime-based visibility for fully capturing string semantics.

The second question (RQ2) addresses how we filter and

*: Daoyuan Wu is the co-first and corresponding author.

organize the training-time string corpus to build a high-quality, semantically meaningful FSS vector database. Naively embedding all raw strings is computationally infeasible and semantically noisy. To address this, we compare a frequency-based filtering method with an LLM-assisted approach that uses GPT-3.5 to estimate string meaningfulness. Our experiments show that LLM-based filtering improves top-1 classification accuracy from 31% to 40%, a relative gain of 29%. Further, we investigate whether intra-family string clustering can improve feature compactness and retrieval quality, laying the groundwork for structurally aware database construction.

The third question (RQ3) considers the test-time phase: given that a malware binary may yield thousands of strings, which subset should be selected as query input to retrieve matching FSS features? We compare random subsampling with a k-means clustering-based method that selects centroid strings using TF-IDF representations [33]. The latter consistently outperforms random sampling, improving classification accuracy by 4% (from 36% to 40%), especially for samples with extensive or diverse string content. This suggests that semantically representative string selection is a critical factor in retrieval-based classification pipelines.

The fourth and final question (RQ4) concerns the choice of inference model. Once the top-k candidate FSS features are retrieved for a test sample, should we score them via vector similarity (i.e., weighted cosine similarity over family-level embeddings) or use an LLM fine-tuned on classification prompts to reason over the retrieved features and infer the family label? Interestingly, our results show that vector-based scoring slightly outperforms the fine-tuned LLM approach (40% vs. 37% top-1 accuracy), despite the latter's semantic capabilities. Error analysis reveals that LLMs struggle with noisy or unparseable input strings, which comprise a significant portion of many binaries. Conversely, vector scoring suffers from higher confusion among semantically similar families. These results suggest complementary strengths, and we outline hybrid inference as a promising direction.

Our experiments are based on a dataset of 4,347 samples across 67 malware families, curated from MalwareBazaar with strict temporal and validation constraints to simulate real-world classification settings. We extract over 25 million raw strings, apply staged filtering and clustering, and build a vector database indexed by meaningful FSS features. Our pipeline achieves 40% top-1 accuracy and 47% top-3 accuracy in the default configuration. More importantly, this pipeline allows us to systematically explore the impact of the four critical design choices described above, laying a foundation for future research into LLM- and RAG-based malware classification.

Open Science and Ethics. We will release the full evaluation code and data after the double-blind review. All experiments were conducted within our own controlled environment.

II. Preliminaries

A. Strings in Binary Samples

Strings embedded in binary executables often reveal valuable semantic information, such as API names, configuration parameters, URLs, or command-line instructions. These strings, even when obfuscated, can provide critical clues for identifying malware behavior and origin. In this work, we primarily leverage *FLOSS* (FireEye Labs Obfuscated String Solver) [4] to extract various categories of strings from malware binaries. FLOSS is a static analysis tool designed to automatically identify and recover obfuscated strings without requiring dynamic execution. Specifically, it detects the following four types of strings:

- **Static Strings**: These are plain ASCII or UTF-16LE strings stored in readable sections (e.g., `.rdata`, `.data`) of the binary. They can be directly extracted without any decoding or emulation. Common examples include hard-coded messages, file paths, or API names.
- **Stack Strings**: These strings are constructed on the stack at runtime, typically through character-wise operations or byte-wise assignments. While not present in the binary as contiguous sequences, stack strings can be recovered through static analysis of instruction patterns that build strings in memory.
- **Tight Strings**: A specific form of stack strings that are tightly encoded and fully decoded within a single function. These are commonly seen in heavily obfuscated malware and are distinguished by their compact encoding and decoding routines.
- **Decoded Strings**: These are strings produced by decryption or decoding routines identified within the binary. FLOSS analyzes known decoding function signatures and instruction sequences to statically recover strings that would otherwise be revealed only at runtime.

B. Prior Works Leveraging String Features for Malware Detection and Classification

Malware detection and classification have long relied on the extraction and engineering of robust features from executable files or their behaviors. Over time, researchers have adopted increasingly diverse features—from static structures like n-grams and headers, to dynamic traces like API call sequences and behavior logs, to domain-specific patterns like DGA-generated URLs. Among these, string features have appeared intermittently, playing various roles depending on task objectives, analysis modality, and modeling preferences.

Table I presents a chronological overview of nine representative studies published between 2001 and 2024, analyzing their respective goals (malware detection vs. family classification), feature engineering strategies, and specifically, how string-based features were handled. The table introduces six dimensions: (1) paper metadata, (2) venue and year, (3) task type, (4) feature type, (5) string role, and (6) how strings are used. This comparison reveals both the historical evolution and methodological divergence in the field.

In early work, Schultz et al. [36] pioneered the use of string features as core discriminative inputs. By leveraging printable strings and n-gram sequences extracted statically from binaries, their system demonstrated strong detection performance, even outperforming traditional antivirus signatures.

TABLE I: Summary of Prior Works Leveraging String Features in Malware Detection and Family Classification

Paper	Venue	Task Type	Feature Type	String Role	How Strings Are Used in the Analysis
Schultz et al.	S&P (2001) [36]	Detection	Strings + n-grams + PE header (static features)	Primary	Data mining using string patterns and n-grams
Rieck et al.	DIMVA (2008) [31]	Detection	API call sequences (dynamic behavior)	Auxiliary	Generate a feature vector for every malware, including the frequency of some specific strings.
Dahl et al.	MALWARE (2015) [34]	Detection	Printable string histograms + binary metadata	Auxiliary	String hash values are used in histogram-based classification; Extract the frequency of selected keywords from disassembled code as feature.
Siddiqui et al.	CODASPY (2016) [37]	Family Classification	Opcode, entropy, and behavior fusion (static)	Auxiliary	Only histograms related to the string length distribution are used.
Huang et al.	DIMVA (2016) [16]	Detection + Family	Dynamic behavior logs (sandbox-based)	Auxiliary	Behavioral signatures are generated by correlating runtime string patterns with system calls.
Dambra et al.	CCS (2023) [11]	Detection + Family	Static + dynamic features (hybrid)	Auxiliary	Static features outperform dynamic; String features had marginal impact.
Aonzo et al.	USENIX (2023) [9]	Family Classification	Human decision features vs. ML features	Auxiliary	Humans and ML shared decision features; strings indirectly referenced.
Wilhelm et al.	RAID (2023) [44]	Detection (evasive)	Behavior summaries (evasion-resilient)	Auxiliary	Behavioral features extracted from string patterns and correlated with API call sequences
Bogdan et al.	RAID (2024) [10]	Detection (DGA)	DNS character-level patterns (network-based)	Auxiliary	Extract statistical features from domains to quantify structural patterns & detect anomalies.

This underscores an important historical moment when strings were treated as primary signals.

Subsequent work such as Rieck et al. [31] and Huang et al. [16] moved away from static indicators and embraced dynamic behavior logs and execution traces. These systems achieved greater generalization and robustness but excluded string features altogether. They reflect a broader trend of relying on runtime behavior as malware authors increasingly obfuscated static content.

In between these poles, some studies have used string features in auxiliary roles. Dahl et al. [34] included hashed printable strings as part of their input feature histograms, showing how string content can be aggregated but abstracted away. Dambra et al. [11], in a comprehensive comparison of static and dynamic features, observed that string-based indicators offered marginal but non-negligible improvements—highlighting their potential when carefully filtered and combined.

Other studies, such as Siddiqui et al. [37] and Wilhelm et al. [44], explicitly avoided string usage in favor of richer statistical or evasive behavior representations. Siddiqui et al. focused on entropy and opcode-level features for family classification, while Wilhelm et al. built classifiers on abstracted behavior summaries, useful against evasive samples.

More recent work by Aonzo et al. [9] explored human-versus-machine feature reliance and revealed that strings, although not used explicitly as input, surfaced in the sandbox reports referenced by both humans and models. This suggests that string-level cues continue to inform decision-making—even if implicitly.

Finally, the Bogdan et al's DGA-focused work [10] illustrates an emerging focus on network-level string-like patterns. By modeling character-level features of dynamically generated domains, it draws on ideas similar to binary string analysis but shifts the focus from executable internals to observable traffic.

C. Rethinking String-based Malware Classification in the New Era of LLMs and RAG

Large Language Models (LLMs), including prominent examples such as the GPT series [30] and LLaMa series [40], [41], represent a significant advancement in natural language processing. Trained on massive, diverse textual corpora, these models exhibit remarkable performance across tasks such as text generation and question answering.

However, general-purpose LLMs—while powerful—often lack domain-specific precision due to the static nature of their pretraining data. To address this, two complementary strategies have emerged: *fine-tuning* [14] and *Retrieval-Augmented Generation* (RAG) [22]. Fine-tuning customizes a pretrained LLM for a particular domain using a smaller, task-specific dataset, improving its ability to produce accurate and context-sensitive responses. In the software engineering domain, fine-tuned models such as CodeLLaMa [32] have proven effective in code generation [24], [43], program repair [15], [38], vulnerability auditing [28], [39], and dynamic analysis [51], [54]. Meanwhile, RAG enhances model outputs by retrieving relevant textual information from an external corpus at inference time [22], enabling LLMs to access up-to-date, domain-specific content [7], [17], [26], [52] without additional training.

These advances motivate a fundamental rethinking of how string-based features—long regarded as static, noisy, or superficial—can be integrated into malware classification pipelines. In traditional systems, string features (e.g., API names, file paths, registry keys, or command-line options) have often been treated as auxiliary or ignored entirely in favor of behavioral or statistical indicators. Yet as demonstrated in early research, strings are rich in semantic content and often highly discriminative, particularly when considered at the family level.

Despite their ubiquity in binary files and their interpretability, string features have often been dismissed in recent malware classification systems. As shown in our review in §II-B, most existing works either exclude string features or embed them

as auxiliary signals within high-dimensional representations. This trend reflects broader concerns about strings being easily obfuscated, sparse, and less reliable than behavior traces or statistical patterns.

However, this prevailing treatment overlooks a key shift: with the emergence of LLMs and retrieval-augmented reasoning, we are now equipped with models that excel at interpreting, comparing, and contextualizing short, unstructured textual data—precisely the form that string artifacts take. LLMs are particularly well-suited to leverage the latent semantics in malware strings, even in the absence of precise structural features or behavioral logs. Moreover, the RAG paradigm provides a natural way to connect an unknown sample's string footprint with a curated corpus of family-specific knowledge, allowing for robust semantic alignment without requiring explicit, symbolic rules or costly model fine-tuning.

This prompts a new perspective: rather than treating strings as pre-classification features to be encoded and abstracted away, they can serve as the *interface layer* between static malware artifacts and LLM-based semantic reasoning. In this view, strings are not merely features—they are the medium through which human-understandable and model-interpretable semantics can be exchanged. Rethinking string-based classification under this paradigm entails (i) curating a high-quality, interpretable set of family-specific strings; (ii) constructing retrieval pipelines that match new samples to relevant historical string contexts; and (iii) leveraging LLMs to reason over the relationships between these artifacts and known malware semantics. This direction combines the transparency of classical static analysis with the adaptability and semantic depth of modern LLMs—offering a novel, interpretable path forward for malware family classification.

III. An Exploratory Study

A. Exploration Based on Family-Specific String Features

Building on this reconceptualization of strings as a semantic interface between malware binaries and LLM-based reasoning, we turn to a systematic exploration of how such strings can be operationalized in practice. Rather than treating string artifacts as low-level, pre-classification features to be abstracted away, we consider them central to a semantic retrieval and reasoning pipeline. Specifically, we investigate how *Family-Specific Strings (FSS)*—interpretable string-level features that are both representative and discriminative of malware families—can be extracted, organized, and utilized throughout the classification process.

We begin by formally defining what constitutes a family-specific string and how it is distinguished from general-purpose string artifacts.

Definition 1 (Family-Specific String)**.** *Let B be a malware binary, and let F denote its corresponding malware family. Let $\mathcal{B}$ be the set of all binary samples that belong to family F, and let $\mathcal{S}$ be the set of all strings decoded from samples in $\mathcal{B}$. Let $\mathcal{F}$ be the set of all known malware families.*

A string S is said to be family-specific *with respect to F if and only if:*

$$S \in \mathcal{S} \quad \textit{and} \quad \nexists F' \in \mathcal{F} \setminus \{F\} \textit{ such that } S \in \mathcal{S}_{F'} \tag{1}$$

where $\mathcal{S}_{F'}$ denotes the set of strings decoded from all binaries associated with family F'.

In other words, a family-specific string is one that occurs within the samples of a given family F but does not appear in any samples belonging to other families.

Overview of the Pipeline. Fig. 1 illustrates the overall pipeline studied in this work. Rather than presenting a fixed system, we decompose the pipeline into four modular components and investigate alternative design choices in each. These components include: (1) string extraction from malware binaries, (2) construction of a searchable vector database of family-level strings, (3) selection of query-time string inputs from test samples, and (4) matching and inference strategies for classification. The goal is to understand how different configurations of these components affect the performance and interpretability of string-centric classification.

This pipeline design supports flexibility in each module, allowing us to isolate and compare key implementation choices. The rest of this section details the research questions and the exploratory experiments conducted in each module.

B. Research Questions

The classification pipeline illustrated in Fig. 1 comprises four key modules: string extraction, vector database construction, query-time string selection, and matching-based inference. Each module presents distinct design choices that influence the pipeline's overall effectiveness and interpretability. Rather than committing to a single architecture, we treat each module as an independent unit of analysis and explore competing strategies through comparative experiments. Our goal is to surface practical insights into how Family-Specific Strings (FSS) can be leveraged for malware family classification in the era of LLMs and retrieval-augmented inference.

To guide this investigation, we articulate four research questions corresponding to the major design modules:

RQ1: Should family-specific strings be extracted purely from static analysis, or can hybrid approaches that incorporate dynamic execution provide better discriminative features?

RQ2: How should we construct the vector database for family-level string matching? (a) How should training strings be compressed: by selecting top strings or using LLM-assisted semantic filtering? (b) Should the strings in the database be further clustered to improve matching quality?

RQ3: Given that test-time samples may contain numerous strings, how can we effectively select observation points (OPs) to serve as the query representation?

RQ4: In the final stage of family classification, should the system rely on similarity-based scoring over retrieved strings, or use fine-tuned LLMs for semantic inference?

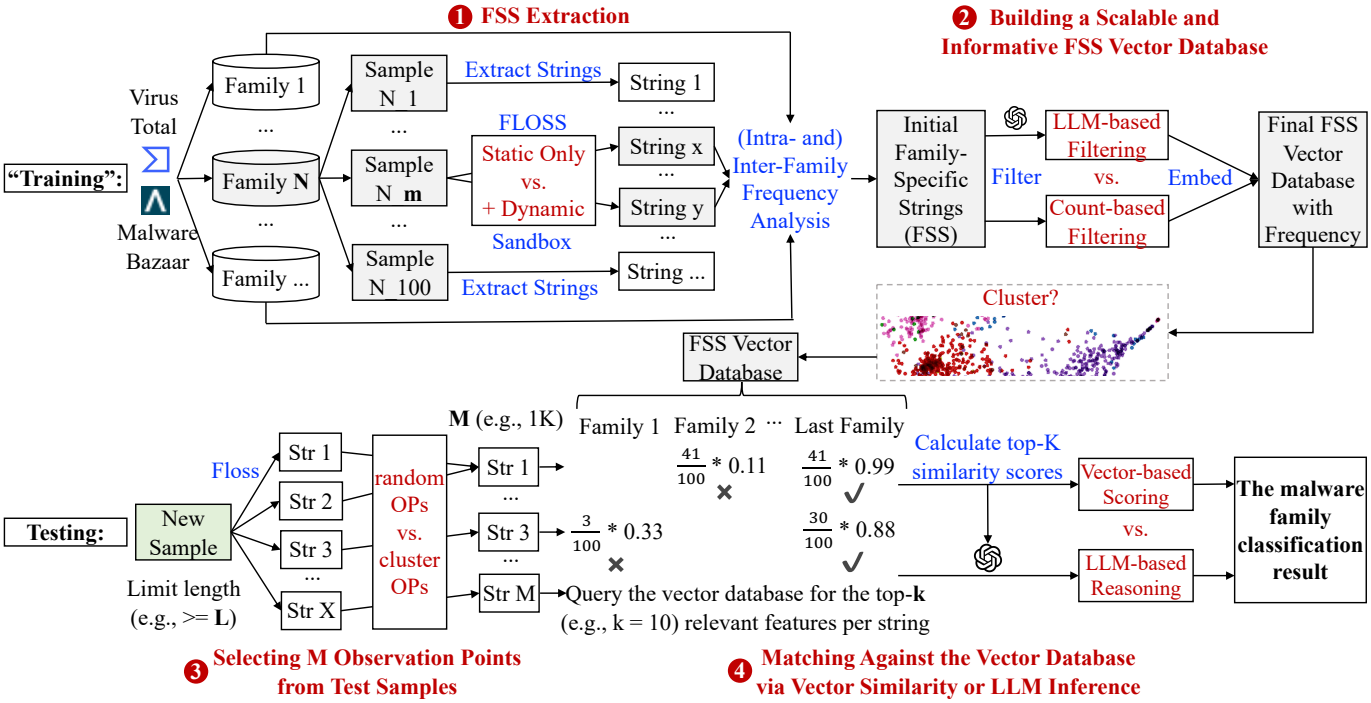

Fig. 1: A high-level overview of our exploratory study.

C. Extracting Family-Specific Strings: Static-Only or Hybrid?

The first component of our exploratory pipeline extracts candidate strings from malware binaries to serve as an initial set of raw FSS features. As illustrated in Fig. 1, this component in the offline "training" phase includes three steps, introduced below; in particular, in the second step, we conduct an ablation study in §V-D to address a foundational question for FSS-based classification: should we rely solely on static analysis, or should we also incorporate dynamic execution to recover a broader and potentially more representative set of string features?

① **Data Collection and Pre-Processing.** As introduced subsequently in §IV, we collect numerous samples for each malware family to represent its variants. Suppose there are N families, as shown in Fig. 1, we collect, process, and obtain m samples for each family, where m could be 100 or 50, depending on the average number of available samples for all families. Specifically, to prioritize samples with more noticeable FSS features for binary string "training," we filter out samples that are *clearly* packed by performing a packing check using Detect It Easy (DIE) [2] and an entropy-based analysis [27].

② **String Extraction from Binaries.** Once we have a set of $N \times m$ samples, we need to extract raw strings from binaries. As mentioned above, there are two design choices: a static-only approach or one that incorporates dynamic execution. The former is scalable and fully automated, while the latter may not be. In §V-C, we compare the effect of the following two strategies for different malware families.

- **Static-only String Extraction**: In the static setting, we utilize FLOSS [4], a widely adopted reverse-engineering tool that extracts four categories of printable strings from binaries: *static strings*, *stack strings*, *tight strings*, and *decoded strings*; see a more detailed introduction in §II-A. FLOSS decodes the embedded strings within each sample, producing a set of raw strings for every malware binary in the family. These strings capture constants, literals, and obfuscated values reconstructed through dataflow analysis. Static methods offer full coverage and reproducibility without requiring execution, but they may fail to expose runtime-generated strings.
- **Incorporating Dynamic String Extraction**: A malware analysis pipeline can also incorporate dynamic execution, although this comes at the cost of scalability and requires manual effort. In this paper, we use Falcon SandBox [1], a commercially supported, feature-rich malware analysis environment. This platform provides real-time API call monitoring with comprehensive coverage of Windows system calls, memory forensics with high string-reconstruction accuracy, and multi-path execution tracing for increased branch coverage. In this hybrid setting, each sample undergoes both FLOSS-based static analysis and Falcon-based dynamic execution. The resulting sets of strings are then unified, with duplicates removed and invalid entries filtered. This setup allows us to empirically assess whether augmenting static FSS with dynamic strings improves family classification performance.

③ **Family Frequency Analysis.** Based on the definition of FSS in Definition 1, we perform a cross-family comparison of all strings extracted by FLOSS to obtain an initial FSS feature set for each malware family, as illustrated in Fig. 1.

In other words, if a string appears in the samples of one family and never appears in any samples of other families, we consider it an FSS. This step filters out common strings that may appear in multiple families, such as general library calls, common system functions, or widely used file names. In addition to inter-family frequency analysis, we also analyze the intra-family frequency of FSS features and sort them by frequency, which will be only used for further filtering when constructing the final FSS vector database in §III-D.

D. Building a Scalable and Informative FSS Vector Database

As illustrated in Fig. 1, the second component involves constructing a scalable and informative vector database from the initial set of raw FSS features obtained in §III-C. It serves as the RAG database that enables query-time retrieval and matching in later stages. However, the raw FSS feature pool is typically large and noisy. For example, in our training corpus of 3,350 samples from 67 families (§IV), the average number of raw FSS strings per family is as high as 334,116 (median: 137,160). Embedding such a large volume of features is computationally expensive and semantically redundant.

To address this challenge, a malware analysis pipeline can employ a two-step refinement process: (1) compressing the raw FSS strings through either LLM-based semantic filtering or top count-based filtering, and (2) optionally organizing them via vector clustering. The effectiveness of these design choices is evaluated in ablation studies, as further discussed in §V-A.

Count-based vs. LLM-based Filtering. A straightforward filtering method is to retrain top-K most frequent FSS features. However, this naive method may ignore FSS features that reflect family-specific features but are not in high count. To address this problem, we also design a filtering method based on semantic content to refine the feature set. This alternative design choice is based on the observation that shorter FSS tend to consist of more meaningless, obfuscated strings (e.g., randomized API names), rather than meaningful, human-readable strings. Leveraging this insight, we employ an LLM-assisted method to determine a length range that primarily captures useful FSS features, thereby enhancing the relevance and interpretability of the embedded features.

LLM-assisted FSS Semantic Analysis. In the alternative LLM-based design choice, we begin by categorizing FSS features based on their string length. Initial analysis revealed that strings longer than a minimum length are more likely to contain meaningful information, such as identifiable keywords, URLs, or other structured patterns. Therefore, we perform semantic analysis of FSS features to *indirectly* determine an appropriate length threshold (i.e., $>= L$). To identify the most valuable FSS features, we employ a prompt-based semantic analysis method using GPT-3.5, as outlined below:

> You are an expert in obfuscated string reading, capable of recognizing any meaning from obfuscated strings. Does each of the following strings have any meaning? Answer in a JSON output only (do not provide any other description), where the key is the string number ID and the value is Yes/No.

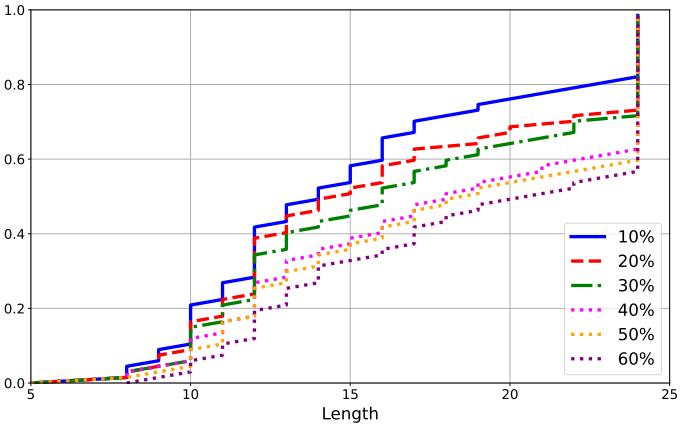

Fig. 2: CDF of meaningful FSS at various percentage levels across different FSS lengths for our training set (§IV).

> 1. string1
> 2. string2
> ...
> 500. string500

For each batch of strings (up to 500 at a time), we input them into the GPT-3.5 model's context window. To mitigate hallucinations and enhance reliability, we use a voting mechanism involving three independent GPT-3.5 agents with different `temperatures`. For each string, if at least two agents agree on "Yes" (i.e., the string is meaningful), it is considered valuable. It is worth noting that our LLM-based semantic analysis is used only to infer the length range where FSS features are most valuable. That said, it does not eliminate any unintelligible strings within the determined length range.

Determining the Optimal FSS Length Threshold. Based on the LLM analysis results, we calculate the percentage of meaningful FSS features for different lengths. Fig. 2 plots the CDF (Cumulative Distribution Function) showing the percentage levels of meaningful FSS across different lengths. We observe that for our training set in §IV, at a length of 13, the minimum percentage of meaningful FSS (i.e., 10%) now covers over 50% of all samples. In other words, for lengths below 13, most samples have a very low percentage of meaningful FSS, often below 10%, making them less suitable for inclusion in the feature database. Therefore, we set the FSS length threshold L to 13 for the samples in our training set. In doing so, we significantly reduce the average (median) number of FSS features per family from 334,116 (137,160) to 33,782 (12,544), achieving an almost 10-fold reduction.

Embedding the Refined FSS Feature Set into Vector Space. With the filtered FSS features through either choice above, we sort them based on the intra-family frequency analysis conducted in §III-C and retain the top 10,000 strings per family. This balances the feature set across families and further reduces the size of our vector database. We then embed this refined set of FSS features into a vector space and store the vector embeddings along with the original plaintext strings, their family labels, and occurrence frequencies. This structure enables efficient approximate nearest-neighbor search, as well as exact match lookup, during test-time matching. Note that

during the embedding of all the refined FSS features, we also record them separately in plaintext to facilitate fast comparison of common FSS features between a testing sample and the training set, as discussed subsequently in §III-E.

Optional Clustering of Embedded Strings. To further reduce redundancy and improve retrieval interpretability, we also explore clustering the embedded FSS vectors using unsupervised algorithms such as *k*-means [20] and DBSCAN [18]. The hypothesis is that grouping semantically similar strings can enhance retrieval stability, reduce query-time noise, and surface core behavioral motifs. We evaluate the impact of clustering in §V-A2 through an ablation analysis.

Overall: Through the above filtering, embedding, optional clustering process, the initial raw FSS features are transformed into a compact and semantically structured knowledge base for test-time retrieval and matching. The trade-offs between filtering strategies and the effect of clustering are key focus areas in our investigation of FSS-based classification.

E. Selecting Observation Points from Test-Time Samples

The third component of our exploratory pipeline concerns how to effectively select *Observation Points (OPs)* from a testing sample to serve as the query representation for family classification. As shown in Fig. 1, each test-time sample typically yields a large number of strings, many of which are noisy, redundant, or irrelevant. Efficiently selecting a meaningful subset of strings is critical for both inference accuracy and computational feasibility.

We explore two competing strategies for selecting OPs: *random subsampling* and *clustering-based selection*. An ablation study in §V-B quantifies the impact of each strategy on classification accuracy, as motivated by RQ3.

Baseline Strategy: Random Subsampling. In the baseline setting, we randomly sample M strings (e.g., $M = 1{,}000$) from the FLOSS-extracted string set of the test sample, after applying the length threshold from §III-D. Moreover, to ensure discriminative string coverage, any common strings shared between the test sample and the FSS vector database (identified through exact string matching) are always included in the OPs. While this random subsampling method is fast and simple, it risks omitting semantically important strings when the total string count is high.

Alternative Strategy: Clustering-based Selection. To address the limitations of random sampling, we introduce a clustering-based method for OP selection. The key intuition is that semantically similar strings often form natural groups, and selecting representatives from these groups can preserve the diversity and informativeness of the entire string set. Specifically, we first convert the string set into vector representations using TF-IDF [33] weighted 3-gram embeddings. Then, we apply the *k*-means algorithm with $k = M$ to cluster the strings. The string closest to the centroid of each cluster is selected as the OP, ensuring coverage across diverse semantic regions of the string space.

Comparative Evaluation. As discussed in §V-B, we evaluate both strategies using a fixed input budget of $M = 1{,}000$ strings per sample. On average, the clustering-based approach improves classification accuracy by over 10% compared to random subsampling, particularly for samples with large and diverse string sets. While clustering-based OP selection improves performance, it also incurs additional computational overhead during preprocessing. The trade-off is justifiable in static analysis settings where string extraction is deterministic and preprocessing is performed offline. Future work may explore more efficient embedding schemes or adaptive clustering thresholds based on string entropy or redundancy levels.

Overall: This module reinforces the need for semantically-aware preprocessing in the test-time pipeline and ensures that test-time representations are both compact and discriminative.

F. Scoring or Reasoning: Matching Against the Vector Database via Vector Similarity or LLM Inference

The final module in our exploratory study concerns how to interpret the relationship between a test-time sample and the family-specific knowledge base constructed from FSS features. As shown in Fig. 1, after retrieving semantically similar features from the vector database, two strategies can be used to perform the final classification: (1) *vector-based scoring*, which ranks families based on similarity-weighted evidence accumulation, and (2) *LLM-based reasoning*, which uses a fine-tuned large language model to infer the sample family based on its top-ranked FSS matches. Our ablation study in §V-C compares these two approaches.

Vector-based Scoring. In this approach, we follow a RAG-like workflow [21], [26] to retrieve and score relevant FSS entries from the database. For each of the M selected observation points (OPs) from the test sample, we query the FSS vector store to retrieve the top-k (e.g., $k = 10$) semantically similar strings. This results in $M \times k$ retrieved strings (or "hints"), each associated with its source family and frequency.

To estimate the relevance of each hint to the test sample, we compute a frequency-normalized similarity score:

$$\text{Score}_{\text{weighted}} = \frac{F_{\text{feature}}}{N_{\text{family}}} \times \text{Sim}(\text{Str}_{\text{OP}}, \text{Str}_{\text{feature}})$$

where F_{feature} is the frequency of the feature string within its source family, N_{family} is the total number of training samples in that family (used for normalization), and $\text{Sim}(\cdot)$ measures semantic similarity in embedding space.

All retrieved hints are ranked by their weighted score, and the top-K matches (e.g., $K = 100$) are used to aggregate family-level scores. The top-ranked family is selected as the predicted label for the sample. This method is lightweight.

LLM-based Reasoning. To enhance semantic interpretation, we further explore an alternative approach where the top-K retrieved features and top-scoring candidate families are passed to a fine-tuned LLM for final inference. As shown in Fig. 3, the model receives two inputs: (1) the top-K retrieved FSS strings and their associated source families, and (2) a ranked list of candidate families based on vector scores. The LLM is trained to output the top-3 candidate families. This is formulated as a structured prompt-based classification

Prompt for Fine-tuning

System:
You are an expert in malware analysis, capable of interpreting the sample type and its family (for in-scope malware) based on the top-K feature strings and top-scoring candidate families.

User:
We have extracted feature strings from a potential malware sample and obtained the top 100 records with the highest scores in an embedding vector database. Below are the most relevant feature strings and their corresponding source families (in brackets) from the given sample (provided as line-by-line input):
{List of top-K string (family) mappings}

Moreover, based on these top feature strings, we have calculated the top potential malware families for the given sample, sorted from highest to lowest likelihood:
{List of top-scoring candidate families}

Based on the above two pieces of information, please determine the sample type (in-scope malware, out-of-scope malware, or benignware) and the top-3 potential families if it is in-scope malware.

Please respond in one of the following three formats (do not include any other descriptions):
The top-3 families are {family1, family2, family3}.

Response:
The top-3 families are {top-1 family, family, family}.

Fig. 3: Prompt format for fine-tuning-based reasoning over retrieved top-K features and candidate families.

task using fine-tuning data derived from validation samples processed through the same pipeline.

We prepare training data for fine-tuning by running the first three modules (§III-C–§III-E) on a held-out validation set. This ensures the model is exposed to realistic top-K feature patterns. We fine-tune a cost-effective LLM (e.g., `gpt-4o-mini`) using this structured dataset, optimizing it to interpret retrieval patterns and make nuanced family-level judgments.

Overall: This module completes the exploratory pipeline by testing the final classification layer's ability to reconcile raw string evidence with semantic family-level inference, offering a direct response to the challenge posed in RQ4.

IV. Experimental Setup

A. Dataset Collection

To support our exploratory study (see Fig. 1), we curated a malware dataset from the MalwareBazaar platform [5], which hosts a large corpus of Windows/UNIX malware samples labeled by multiple engines. Our dataset construction emulates a real-world scenario where analysts must detect newly emerging malware using a historical repository of known families.

TABLE II: Summary of datasets used.

Dataset	Samples	Families
Training (Vector DB)	3,350	67
Fine-tuning Training Set	670	67
Testing Set (2024 samples)	327	67
Total	4,347	67

We collected and curated a subset of 4,347 labeled samples spanning 67 distinct families for training, validation, and testing. We validated family labels with VirusTotal to remove inconsistently tagged samples. For each malware family, we selected 60 pre-2023 samples (50 for the database and 10 for fine-tuning) and 5 near-January 2024 samples for testing. Some families (e.g., DanaBot, Meterpreter, AgentTesla, Metasploit) had fewer than 5 testable samples due to extraction or obfuscation issues. As summarized in Table II, the dataset is divided into:

- **Training Set**: 3,350 samples (50 per family) collected before 2023, used to construct the FSS vector database.
- **Fine-tuning Set**: 670 samples (10 per family), used to fine-tune the LLM for semantic inference (see §III-F).
- **Testing Set**: 327 samples (up to 5 per family) from early 2024, used to evaluate generalization to recent threats.

B. Environment and Parameter Configurations

Our experiments were conducted on a workstation with 128GB RAM, 32-core CPU, and NVIDIA A100 GPU. The implementation uses Python with `faiss` for vector indexing and OpenAI APIs for embedding and fine-tuning.

Embedding and Query Pipeline. We use the `text-embedding-ada-002` model to convert each FSS string into a 1,536-dimensional embedding vector [3]. During inference, each observation point (OP) from a test-time sample is used to query the FSS vector database. The system retrieves the top-$k = 10$ semantically similar entries per OP, based on cosine similarity, forming the retrieval foundation for both scoring-based and LLM-based reasoning.

LLM Fine-tuning. For semantic inference, we fine-tune the `gpt-4o-mini-2024-07-18` model using our curated prompt-response pairs. Each prompt encodes the top-$K = 100$ retrieved FSS strings and their associated source families, along with a ranked list of candidate families (see Fig. 3). The fine-tuning set of 670 samples is processed using the same pipeline described in §III-C–§III-E, ensuring the model learns to interpret retrieved features.

Default Parameters. Unless otherwise specified, we set $M = 1{,}000$ as the number of OP strings selected from each test sample, $k = 10$ as the number of retrieved FSS strings per OP, and $K = 100$ as the total number of top-ranked FSS features used for scoring or LLM-based classification.

C. How Strings Evolve in the Training Set

To better understand how our exploratory pipeline contributes to refining string-based features, we trace the evolution

of string sets throughout the offline training phase. Specifically, we analyze how the average number of string features per family changes as we successively apply the design choices discussed in §III-C and §III-D. Our experiments were conducted on the 67 malware families used for training.

Step 1: Raw String Extraction. We begin by extracting all printable and decoded strings from each binary using FLOSS, resulting in an initial set of raw strings per sample. On average, this process yields 377,034 strings per family, reflecting the large volume and typical redundancy of FLOSS outputs.

Step 2: Family-Specific Strings. Next, we apply inter-family filtering to identify FSS strings that are unique to each family (§III-C). This step removes strings common across multiple families (e.g., general system API names), reducing the average string count to 334,116 per family.

Step 3: LLM-Based Filtering. To eliminate short and often unintelligible strings, we apply the minimum length threshold ($L = 13$), derived via LLM-assisted semantic analysis (§III-D). This step filters out a large fraction of noisy strings and reduces the average string count to 33,782 per family.

Step 4: Removing Recursive Strings. As an implementation refinement, we remove recursive strings—those composed of repeating self-patterns (e.g., "ABABABAB" or "XYZXYZXYZ")—which typically result from FLOSS decoding loops or malformed encodings. This filtering step eliminates another small subset of strings, bringing the average count down slightly to 33,283.

Step 5: Frequency-Based Compression. Finally, to ensure scalability and balance across families, we retain only the top 10,000 FSS features per family based on intra-family frequency (§III-D). This final pruning step further reduces the average to 7,042 strings per family—an 53-fold reduction from the original FLOSS output—while preserving the most representative and discriminative features.

Through this sequential filtering pipeline, our system transforms a noisy, oversized raw string set into a compact, interpretable, and semantically meaningful feature corpus. This refined dataset forms the backbone of our FSS vector database and directly impacts downstream classification performance.

V. Results and Analysis

We now evaluate the classification performance of our exploratory pipeline across 67 malware families using the curated testing dataset introduced in §IV. Unless otherwise specified, all reported results follow the *Normal Classification* configuration, which incorporates the following default design choices:

- Static-only string extraction using FLOSS (i.e., dynamic analysis is excluded) in RQ1.
- LLM-based filtering of FSS without additional clustering of embedded vectors in RQ2.
- Clustering-based selection of observation points in RQ3.
- Vector-based scoring for family prediction in RQ4.

This configuration represents a fully automated pipeline optimized for speed and scalability. Two modules—dynamic string extraction (RQ1) and vector clustering (RQ2)—are omitted in the default pipeline due to their dependency on manual analysis or non-trivial orchestration. These components are revisited in later subsections through ablation studies to assess their potential impact.

Table III summarizes the per-family and average performance under various configurations. In the default setting (Normal Classification), our pipeline achieves an average top-1 accuracy of 40% across all test samples, indicating that in 40% of cases, the correct malware family is ranked first. When expanding the evaluation window to include the top-2 and top-3 ranked predictions, the average accuracy increases to 45% and 47%, respectively. This incremental improvement (+5%, +2%) suggests that the correct family could be ranked among the top candidates, even if not in top-1, underscoring the semantic value of the retrieved string features.

To understand the effect of each module in the pipeline, we organize the subsequent subsections around the four RQs defined in §III-B. The presentation proceeds in the following order: starting with §V-A (RQ2), which examines how to build an effective FSS vector database; followed by §V-B (RQ3), which compares different strategies for selecting observation points; then §V-C (RQ4), which contrasts vector-based scoring with LLM-based reasoning for final classification; and finally, §V-D (RQ1), which investigates the effect of incorporating dynamic execution into the string extraction phase.

Together, these analyses shed light on the strengths and limitations of the different modules in our pipeline, offering insights into how string-based classification can be enhanced through modular design choices in the era of LLMs and RAG.

A. RQ2: How to Build the FSS Vector Database?

1) RQ2.1: LLM-based Filtering vs. Count-based Filtering: We evaluate the effectiveness of LLM-based length filtering by conducting comparative experiments across 67 malware families. Two experimental conditions are configured:

- **Top Count-based Filtering (Baseline)**: Retain the top 10,000 most frequent strings per family, without applying any semantic or length-based filtering. This simple heuristic yields a top-1 classification accuracy of 31%, as reported in Table III.
- **LLM-based Length Filtering (Default)**: Apply a length threshold (e.g., $L \geq 13$) derived from the LLM-driven analysis described in §III-D, which aims to exclude semantically uninformative short strings. This approach improves top-1 accuracy to 40.3%, i.e., 0.40 in Table III.

As shown in Table III, LLM-based filtering leads to a 29% relative improvement in top-1 accuracy (from 31% to 40%), and similarly improves the top-2 (from 36% to 45%) and top-3 (from 40% to 47%) metrics. These results validate the hypothesis that short or obfuscated strings introduce semantic noise detrimental to classification.

To further understand the semantic improvement brought by LLM-based filtering, we analyze changes in intra-family string similarity. As shown in Fig. 4a, 40% of families (27 out of 67) exhibit improved semantic coherence, reflected by positive Δ values. The mean improvement across these families is

TABLE III: Classification accuracy under different situations for samples in different malware families.

Family	Normal Classification			Use Count Filtering–RQ2			Use Random OPs–RQ3			Use LLM Reasoning–RQ4		
	top1	top2	top3	top1	top2	top3	top1	top2	top3	top1	top2	top3
Adware.Generic	0.00	0.00	0.00	0.40	0.40	0.40	0.00	0.00	0.00	0.00	0.40	0.40
AgentTesla	0.20	0.40	0.40	0.00	0.20	0.20	0.20	0.40	0.40	0.20	0.20	0.40
Amadey	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
AsyncRAT	0.80	0.80	0.80	0.40	0.80	0.80	0.60	0.80	0.80	0.80	0.80	0.80
AveMariaRAT	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
AZORult	0.20	0.40	0.40	0.40	0.40	0.40	0.20	0.40	0.40	0.20	0.40	0.40
BazaLoader	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
BitRAT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
BumbleBee	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CobaltStrike	0.20	0.20	0.60	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
CoinMiner	0.20	0.40	0.40	0.00	0.20	0.20	0.20	0.40	0.40	0.20	0.20	0.20
DanaBot	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DarkCloud	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
DarkComet	1.00	1.00	1.00	0.00	0.00	0.00	1.00	1.00	1.00	0.00	0.00	0.00
DarkGate	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
DCRat	0.40	0.40	0.40	0.20	0.20	0.20	0.60	0.60	0.60	0.20	0.20	0.20
Fabookie	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Formbook	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Gafgyt	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Glupteba	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
GootLoader	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Guildma	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.20
GuLoader	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.00	0.00	0.00
Havoc	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80
Heodo	0.20	0.20	0.20	0.00	0.00	0.00	0.20	0.20	0.20	0.00	0.00	0.00
IcedID	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Kaiji	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Kutaki	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Lazarus	0.40	0.40	0.40	0.60	0.60	0.60	0.20	0.60	0.60	0.40	0.60	0.60
LimeRAT	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Loki	0.60	0.60	0.60	0.00	0.00	0.00	0.60	0.60	0.60	0.00	0.00	0.00
MassLogger	0.40	0.40	0.40	0.20	0.20	0.20	0.20	0.20	0.20	0.00	0.20	0.20
Metasploit	0.25	0.25	0.25	0.25	0.25	0.75	0.25	0.25	0.25	0.25	0.25	0.50
Meterpreter	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Mirai	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
NanoCore	0.00	0.40	0.40	0.20	0.20	0.40	0.20	0.40	0.40	0.60	0.60	0.60
Neshta	0.40	0.40	0.40	0.00	0.20	0.40	0.40	0.40	0.40	0.20	0.20	0.40
NetSupport	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.40	0.40	0.40
Nitol	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40
njrat	1.00	1.00	1.00	0.20	0.20	1.00	0.20	1.00	1.00	1.00	1.00	1.00
Phorpiex	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.80	0.80	0.80
Pikabot	0.00	0.20	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Pony	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
PrivateLoader	0.00	0.00	0.00	0.00	0.20	0.20	0.00	0.00	0.00	0.20	0.20	0.20
PureCrypter	0.20	0.60	0.60	0.00	0.20	0.20	0.20	0.20	0.20	0.40	0.40	0.40
QuasarRAT	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
RecordBreaker	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40
RevengeRAT	0.60	0.60	0.60	0.40	0.40	0.40	0.40	0.60	0.60	0.40	0.40	0.40
Rhadamanthys	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
RustyStealer	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
SnakeKeylogger	0.20	0.20	0.20	0.00	0.00	0.00	0.00	0.00	0.20	0.00	0.00	0.00
Stealc	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Stop	0.40	0.60	0.60	0.40	1.00	1.00	0.80	1.00	1.00	0.40	0.80	1.00
StormKitty	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
StrelaStealer	0.40	0.60	0.60	0.00	0.20	0.20	0.00	0.20	0.60	0.20	0.20	0.20
STRRAT	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
SystemBC	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60	0.60
TeamBot	0.20	0.20	0.60	0.00	0.00	0.40	0.00	0.00	0.00	0.20	0.20	0.60
Tofsee	0.00	0.20	0.80	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Tsunami	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Vidar	0.20	0.40	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Vjw0rm	0.40	0.40	0.40	0.00	0.40	0.40	0.00	0.40	0.40	0.40	0.40	0.40
WSHRAT	0.60	0.80	0.80	0.60	0.60	0.60	0.60	0.80	0.80	0.40	0.60	0.60
XWorm	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
YellowCockatoo	1.00	1.00	1.00	0.80	0.80	0.80	1.00	1.00	1.00	0.80	0.80	0.80
ZeuS	0.60	1.00	1.00	0.00	0.20	0.80	0.00	0.20	0.20	0.60	0.60	0.60
zgRAT	0.00	0.20	0.20	0.00	0.00	0.20	0.00	0.20	0.20	0.20	0.20	0.20
Average Sample	**0.40**	**0.45**	**0.47**	**0.31**	**0.36**	**0.40**	**0.36**	**0.42**	**0.43**	**0.37**	**0.40**	**0.42**

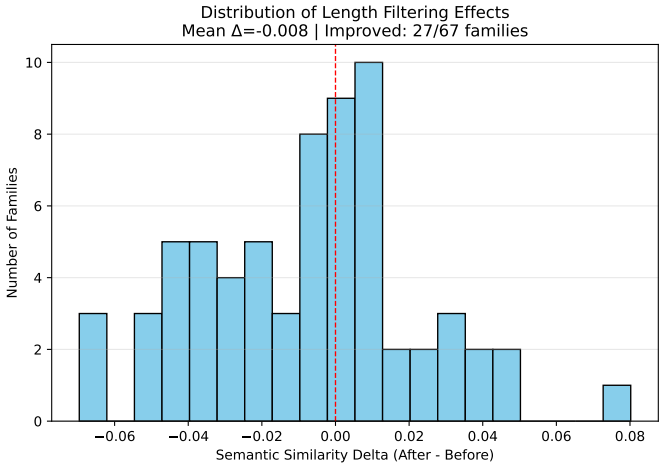

(a) Similarity delta value distribution across families.

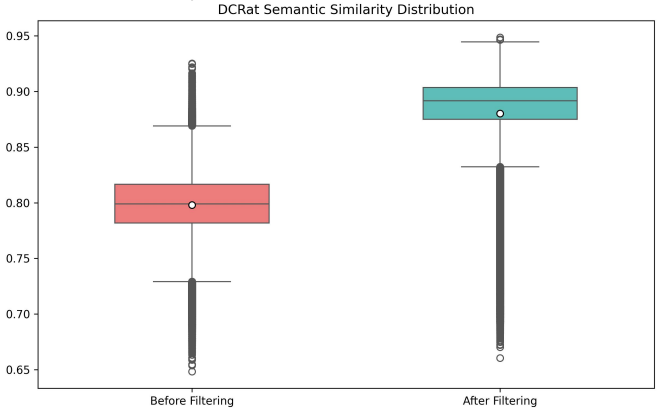

(b) Case study of the DCRat family.

Fig. 4: Impact of LLM-based filtering on semantic similarity.

TABLE IV: Family-wise Changes with LLM-based Filtering.

Family	Before	After	Δ
DCRat	0.720	0.800	+0.080
CoinMiner	0.734	0.782	+0.048
RecordBreaker	0.792	0.723	−0.069
CobaltStrike	0.785	0.723	−0.062

$+0.021$, suggesting that LLM filtering tends to remove noisy or low-information strings that dilute family-specific patterns.

Among the families, DCRat demonstrates the most pronounced benefit, with a delta of $+0.080$ after filtering, as illustrated in Fig. 4b. While the overall improvements are encouraging, the effectiveness of LLM-based filtering is not uniform across all malware families. We observe a substantial degree of variance in delta values, with a standard deviation of $\sigma_\Delta = 0.030$, indicating that *the semantic effect of filtering is highly dependent on malware family characteristics*.

As observed in both Table IV and Fig. 4, the gain in semantic clarity varies significantly depending on the structural characteristics of the family. In particular, we discover that filtering is more effective for malware families whose semantic content tends to correlate with string length. These include families such as DCRat and CoinMiner, where command-and-control (C2) payloads or embedded commands often follow fixed-length patterns. In contrast, for families such as RecordBreaker and CobaltStrike—where obfuscation is more common and semantically meaningful strings may be short or deliberately fragmented—length-based filtering may inadvertently remove useful features.

TABLE V: Cluster Types vs. Classification Performance.

Cluster Type	Jaccard Similarity	Accuracy	Prevalence
Clustered	>0.3	0.87 ± 0.05	61.8%
Non-Clustered	0	0.31 ± 0.12	38.2%

These results support our hypothesis that semantic filtering, guided by LLMs, is especially beneficial under two key conditions. First, when malware families exhibit consistent length-encoded behavior—such as well-formed C2 commands or templated configuration strings—filtering out short, low-entropy strings helps retain high-value patterns. Second, when the length threshold (e.g., $L = 13$) aligns with the operational norms of that family, the filtering process strikes a favorable balance between noise reduction and information preservation.

Nonetheless, the fixed nature of our thresholding strategy introduces limitations. Our current pipeline adopts a global threshold of $L = 13$ for all families, based on the aggregated LLM analysis results. However, the family-level heterogeneity observed suggests that this one-size-fits-all approach may be suboptimal. Future work should explore adaptive or family-specific length thresholds that can better accommodate the diversity in string structures and improve the generalization of semantic filtering across different malware types.

2) RQ2.2: How Do Clusters Affect Prediction?: Our experiments uncover a significant performance gap across malware families, with roughly 40% achieving high classification accuracy (above 0.8), while the remaining 60% show substantially weaker results. Upon closer inspection, this variation is strongly associated with the presence or absence of well-formed file clusters—subsets of samples within a family that shared common frequent strings. These clusters emerge naturally through string co-occurrence patterns: for instance, if Sample A contained {String X, String Y} and Sample B contained {String X, String Z}, both samples would be grouped due to their shared String X. This clustering behavior suggests structural or behavioral similarities within certain variants of a family.

To better characterize this phenomenon, we group samples into two categories. First, *Clustered Families* (61.8% of cases), which exhibit overlapping string patterns and yield an average classification accuracy of 0.87. Second, *Non-Clustered Samples* (38.2%), which lack discernible intra-family string consistency and see their accuracy drop to 0.31. Table V summarizes these differences in Jaccard similarity, accuracy, and dataset prevalence.

To understand why some samples defy clustering, we manually analyze 50 randomly selected non-clustered samples. Two root causes emerge. The majority (64%) are affected by *obfuscation artifacts*, including runtime packers, encryption, and custom string encodings that masked meaningful features. The remaining 36% are identified as *genetic variants*—samples

TABLE VI: Performance comparison of sampling methods (mean ± 95% confidence interval).

Method	Accuracy	Relative Improvement
Random sampling	36%	–
K-means clustering	40%	+11.1%

with substantially different codebases despite sharing a nominal family label. These outliers may reflect version drifts or parallel development tracks within the same malware lineage.

These findings carry several implications. First, high performance among clustered families reinforces the value of string-based features when variants retain shared functional components. Second, the difficulty of classifying non-clustered samples illustrates fundamental limits of static string analysis under heavy obfuscation. Third, cluster detection itself could be leveraged as a preprocessing heuristic to triage samples: well-clustered inputs can proceed via lightweight string-matching inference, while non-clustered ones can be routed to more advanced methods such as dynamic behavior profiling or graph-based structural comparison.

These observations suggest a need for cluster-aware classification pipelines. By dynamically identifying the clustering structure of test-time samples, the pipeline can selectively deploy efficient, string-based techniques for homogenous groups and fallback strategies for obfuscated or divergent variants.

B. RQ3: Random Subsampling vs. Clustering-based Selection

In this RQ, we investigate the effect of different observation point (OP) selection strategies on classification performance. Since each malware sample may yield thousands of decoded strings—even after filtering—selecting a representative and informative subset for querying the FSS vector database is critical. To address this, we compare a baseline random sampling method with a clustering-based selection strategy.

- **Random Subsampling (Baseline)**: Sample 1,000 strings uniformly at random from each test sample's post-filtered string set. While simple and computationally efficient, this strategy often fails to capture semantically diverse strings, particularly in large string sets.
- **Clustering-based Selection (Default)**: Perform k-means clustering over the sample's string set using TF-IDF features with character 3-grams and cosine similarity. From each cluster, the string nearest to the centroid is selected. This ensures semantic diversity by covering multiple regions of the string space, while maintaining a consistent OP count of $M = 1000$.

As shown in Table VI, clustering-based OP selection achieves a top-1 accuracy of 40%, compared to 36% using random sampling—an 11.1% relative improvement. The advantage is particularly evident for samples with more than 10,000 decoded strings, where redundancy and noise are more prevalent. The centroid-based selection improves the representativeness of the OPs, which in turn leads to better alignment with family-specific features in the vector database.

TABLE VII: Classification Accuracy Comparison Between LLM- and Vector-based Approaches.

Approach	Accuracy	Relative Improvement
LLM-based Reasoning	37%	–
Vector-based Scoring	40%	+8.1%

TABLE VIII: Error Type Distribution Across Approaches.

Error Type	Vector-based	LLM-based
Unparseable Strings	37%	55%
Similar Family Confusion	25%	7%

C. RQ4: Vector-based Scoring vs. LLM-based Reasoning

In this RQ, we investigate the effectiveness of two competing strategies for malware family classification: (1) a LLM-based reasoning method (baseline) that employs a fine-tuned LLM to infer malware families from retrieved evidence, and (2) a vector-based similarity scoring method (default) relying on vector similarity and frequency-weighted scoring.

Table VII presents the classification performance under both approaches. The vector-based scoring method achieves a top-1 accuracy of 40%, while the LLM-based reasoning approach yields a slightly lower accuracy of 37%. This modest 3% gap corresponds to an 8.1% relative improvement in favor of vector-based scoring, suggesting that despite the semantic capabilities of LLMs, their advantages do not materialize clearly under current conditions.

To understand the source of performance differences, we performed a detailed error breakdown, as shown in Table VIII. The most common source of error in both approaches is unparseable input—strings that are either nonsensical or the result of obfuscation. These strings account for 37% of misclassifications in vector-based scoring and 55% in the LLM-based reasoning. Conversely, vector-based scoring is more prone to confusing samples from semantically similar families (25% vs. 7% in LLMs), possibly due to its reliance on surface-level pattern matching without contextual understanding.

These results suggest that LLMs exhibit better semantic generalization in distinguishing closely related families, but are more sensitive to noisy or malformed inputs. This aligns with expectations: LLM-based reasoning benefits from contextual understanding when inputs are meaningful, but suffers when input quality is low. On the other hand, vector-based scoring—though more brittle semantically—handles noisy or shallow patterns more robustly due to its statistical nature.

Given these observations, we draw three key conclusions:

1) In environments where input strings are heavily obfuscated or contain low semantic content, vector-based scoring remains competitive and more stable than LLMs.
2) LLMs may offer benefits in handling ambiguous or borderline cases, particularly when the distinction between malware families is subtle and pattern-based similarity is insufficient.
3) A promising future direction lies in hybrid models that combine both strategies—e.g., using vector scoring to prune irrelevant candidates and invoking LLM-based inference only on the top-ranked clusters.

TABLE IX: Comparison of meaningful vs. obfuscated strings.

Meaningful Strings	Obfuscated Strings
`CryptEncrypt`, `Microsoft\Proof\hyph32.dll`, `https://www.ic.ncsoft.com`	`Xj3$kP*9@m`, `aGVsbG8g8J+Yig==`
Average Rate: 31.2%	Average Rate: 68.8%

D. RQ1: Static-only vs. Hybrid String Extraction

This RQ examines the extent to which dynamic execution improves the quality of extracted strings and, in turn, malware family classification performance. Our default pipeline uses FLOSS for static-only string extraction due to its scalability and automation. However, as malware increasingly adopts obfuscation techniques and runtime packing, static methods alone may fail to uncover semantically useful information.

Our initial, preliminary evaluation revealed that around one third of samples exhibited poor classification accuracy when relying solely on static extraction. To investigate this, we manually analyze the string artifacts and observe that, even after entropy filtering and regular expression pruning, many static strings remained obfuscated or meaningless. Table IX presents representative examples contrasting meaningful strings (e.g., API names or URLs) with obfuscated strings (e.g., base64 junk or byte repetitions). On average, only 31.2% of extracted strings are semantically interpretable, while 68.8% are either garbled or machine-generated filler.

As a comparative experiment, we introduce a hybrid string extraction method that supplements static outputs with dynamic traces collected from Falcon Sandbox. Falcon provides runtime visibility of system calls, decrypted payloads, and memory-resident strings—many of which are inaccessible statically. We select five representative families with poor static-only classification (Loki, Stop, NanoCore, Formbook, and DarkGate) and rerun the pipeline with dynamic augmentation.

The results are striking. On average, we observe a 63.2% increase in the number of semantically valid strings, reflecting Falcon's ability to capture runtime-generated content such as decrypted payloads and system interactions. Additionally, the proportion of garbage or obfuscated strings drop by 41.5%, indicating that dynamic execution is effective at bypassing static noise introduced by packing and encoding. Perhaps most notably, the identification of system-level API calls improve by a factor of 2.8×, dramatically enhancing the visibility of behavioral indicators crucial for malware family classification.

As shown in Table X, these enhancements translate into clear gains in classification performance. For instance, the `Stop` ransomware family sees its accuracy jump from 0.2 (static-only) to 1.0 (hybrid), driven by better visibility into encryption routines, file system interactions, and C2 URLs—all of which were hidden from static inspection. In contrast, families like `DarkGate` show no improvement, suggesting that even dynamic execution may be insufficient when anti-analysis techniques or stealth behaviors are involved.

These results confirm the hypothesis that dynamic string extraction offers value for certain malware families. While not universally beneficial, hybrid string extraction could be considered a critical capability. In future work, we plan to explore partial automation of dynamic execution (e.g., using unpacking stubs or emulated environments) to expand hybrid coverage without sacrificing scalability.

TABLE X: Classification accuracy before and after dynamic execution for certain malware families.

Family	Static-only	Hybrid (Static+Dynamic)
Loki	0.60	0.60
Stop	0.20	1.00
NanoCore	0.20	0.40
Formbook	0.00	0.20
DarkGate	0.00	0.00
Average	0.20	0.44

VI. Discussion

Our study reveals both the promise and limitations of string-based malware classification within the context of LLMs and retrieval-augmented pipelines. One further extension is the importance of *dynamic length filtering*, tailored specifically to family characteristics. While we adopted a fixed threshold ($L = 13$) based on aggregated LLM feedback (§V-A1), our results show that families such as `DCRat` and `CoinMiner` benefit from this filtering, whereas others like `CobaltStrike` and `RecordBreaker` are negatively impacted. This highlights the need for adaptive, family-aware filtering to preserve meaningful short strings that might otherwise be discarded.

We also identify limitations in *LLM fine-tuning* on abstract string data. Although LLMs excel at code and text-related tasks, their effectiveness here is limited (§V-C) due to two main factors: (1) retrieved FSS strings often lack clear semantics—particularly in packed or obfuscated binaries—making it challenging for LLMs to establish reliable associations; and (2) obfuscation artifacts in training data introduce misleading patterns, reducing generalization. These findings suggest that combining RAG-based retrieval with more structured semantic filtering may provide a more robust alternative.

VII. Related Work

Beyond the summary of string-based classification methods surveyed in §II-B, our work is situated at the intersection of malware analysis and the emerging use of LLMs in cybersecurity. LLMs have recently demonstrated strong capabilities across a variety of security tasks: PentestGPT [12] and CHATAFL [29] applied LLMs to penetration testing and protocol fuzzing, while other systems have explored fuzzing of libraries [13], [47], program repair [52], [53], and binary analysis [25], [45]. Code-centric applications such as LLift [23] and GPTScan [39] further highlight LLMs' strengths in static code understanding. However, existing research has largely overlooked the integration of LLMs with string-level artifacts in malware binaries—a space traditionally seen as noisy. Our work addresses this gap by demonstrating that, with appropriate filtering and semantic grounding, string features can serve as an effective interface between binary artifacts and LLM-based reasoning, offering a lightweight, interpretable alternative to deeper instrumentation or opaque embeddings.

VIII. Conclusion

This study revisited string-based malware family classification in the modern context of LLMs and retrieval-augmented architectures. By introducing the concept of Family-Specific Strings (FSS) and designing a modular pipeline centered on their extraction, organization, and inference, we demonstrated that string artifacts can, when properly curated, yield semantically meaningful and interpretable signals for malware classification. Our exploratory analysis across four key design stages revealed that hybrid string extraction improves feature coverage for obfuscated malware, LLM-based filtering enhances semantic consistency, clustering-based observation point selection improves representation quality, and vector-based scoring offers robust and efficient inference.

References

[1] "CrowdStrike Falcon Sandbox." [Online]. Available: https://www.crowdstrike.co.uk/products/threat-intelligence/falcon-sandbox-malware-analysis/

[2] "Detect It Easy." [Online]. Available: https://github.com/horsicq/Detect-It-Easy/

[3] "Embedding Model: text-embedding-ada-002." [Online]. Available: https://openai.com/index/new-and-improved-embedding-model/

[4] "FLOSS: FLARE Obfuscated String Solver." [Online]. Available: https://github.com/mandiant/flare-floss

[5] "MalwareBazaar." [Online]. Available: https://bazaar.abuse.ch/

[6] "VirusTotal." [Online]. Available: https://www.virustotal.com/

[7] "Llm4vuln: Large language models for vulnerability detection," Jan. 2024. [Online]. Available: https://sites.google.com/view/llm4vuln/home

[8] H. Aghakhani, F. Gritti, F. Mecca, M. Lindorfer, S. Ortolani, D. Balzarotti, G. Vigna, and C. Kruegel, "When malware is packin'heat; limits of machine learning classifiers based on static analysis features," in *Network and Distributed System Security Symposium*. Internet Society, 2020.

[9] S. Aonzo and collaborators, "Humans vs. machines in malware classification," in *USENIX Security Symposium*, 2023.

[10] R. Consortium, "Down to earth! guidelines for dga-based malware detection," in *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*, 2024.

[11] S. Dambra, Y. Han, S. Aonzo, P. Kotzias, A. Vitale, J. Caballero, D. Balzarotti, and L. Bilge, "Decoding the secrets of machine learning in malware classification: A deep dive into datasets, feature extraction, and model performance," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 60–74.

[12] G. Deng, Y. Liu, V. Mayoral-Vilches, P. Liu, Y. Li, Y. Xu, T. Zhang, Y. Liu, M. Pinzger, and S. Rass, "Pentestgpt: An llm-empowered automatic penetration testing tool," *arXiv preprint arXiv:2308.06782*, 2023.

[13] Y. Deng, C. S. Xia, H. Peng, C. Yang, and L. Zhang, "Large language models are zero-shot fuzzers: Fuzzing deep-learning libraries via large language models," in *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. Seattle WA USA: ACM, Jul. 2023, p. 423–435. [Online]. Available: https://dl.acm.org/doi/10.1145/3597926.3598067

[14] N. Ding, Y. Qin, G. Yang, F. Wei, Z. Yang, Y. Su, S. Hu, Y. Chen, C.-M. Chan, W. Chen *et al.*, "Parameter-efficient fine-tuning of large-scale pre-trained language models," *Nature Machine Intelligence*, vol. 5, no. 3, pp. 220–235, 2023.

[15] K. Huang, J. Zhang, X. Bao, X. Wang, and Y. Liu, "Comprehensive fine-tuning large language models of code for automated program repair," *IEEE Transactions on Software Engineering*, 2025.

[16] W. Huang and J. W. Stokes, "Mtnet: A multi-task neural network for dynamic malware classification," in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2016, pp. 399–418.

[17] Z. Ji, D. Wu, W. Jiang, P. Ma, Z. Li, and S. Wang, "Measuring and augmenting large language models for solving capture-the-flag challenges," in *ACM CCS*, 2025.

[18] K. Khan, S. U. Rehman, K. Aziz, S. Fong, and S. Sarasvady, "Dbscan: Past, present and future," in *The fifth international conference on the applications of digital information and web technologies (ICADIWT 2014)*. IEEE, 2014, pp. 232–238.

[19] D. Kirat and G. Vigna, "Malgene: Automatic extraction of malware analysis evasion signature," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 2015, pp. 769–780.

[20] K. Krishna and M. N. Murty, "Genetic k-means algorithm," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 29, no. 3, pp. 433–439, 1999.

[21] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.

[22] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, "Retrieval-augmented generation for knowledge-intensive nlp tasks," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, p. 9459–9474. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf

[23] H. Li, Y. Hao, Y. Zhai, and Z. Qian, "Enhancing static analysis for practical bug detection: An llm-integrated approach," *Proceedings of the ACM on Programming Languages*, vol. 8, no. OOPSLA1, pp. 474–499, 2024.

[24] Z. Li, C. Wang, P. Ma, C. Liu, S. Wang, D. Wu, C. Gao, and Y. Liu, "On extracting specialized code abilities from large language models: A feasibility study," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.

[25] P. Liu, C. Sun, Y. Zheng, X. Feng, C. Qin, Y. Wang, Z. Li, and L. Sun, "Harnessing the power of llm to support binary taint analysis," no. arXiv:2310.08275, Oct. 2023, arXiv:2310.08275 [cs]. [Online]. Available: http://arxiv.org/abs/2310.08275

[26] Y. Liu, Y. Xue, D. Wu, Y. Sun, Y. Li, M. Shi, and Y. Liu, "PropertyGPT: LLM-driven Formal Verification of Smart Contracts through Retrieval-Augmented Property Generation," in *Proc. ISOC NDSS*, 2025.

[27] R. Lyda and J. Hamrock, "Using entropy analysis to find encrypted and packed malware," *IEEE Security & Privacy*, vol. 5, no. 2, pp. 40–45, 2007.

[28] W. Ma, D. Wu, Y. Sun, T. Wang, S. Liu, J. Zhang, Y. Xue, and Y. Liu, "Combining fine-tuning and llm-based agents for intuitive smart contract auditing with justifications," in *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. IEEE Computer Society, 2024, pp. 330–342.

[29] R. Meng, M. Mirchev, M. Bohme, and A. Roychoudhury, "Large language model guided protocol fuzzing," in *Proceedings of the Symposium on Network and Distributed System Security 2024*.

[30] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe, "Training language models to follow instructions with human feedback," Mar. 2022, arXiv:2203.02155 [cs]. [Online]. Available: http://arxiv.org/abs/2203.02155

[31] K. Rieck, T. Holz, C. Willems, P. Düssel, and P. Laskov, "Learning and classification of malware behavior," in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2008, pp. 108–125.

[32] B. Rozière, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, T. Remez, J. Rapin, A. Kozhevnikov, I. Evtimov, J. Bitton, M. Bhatt, C. C. Ferrer, A. Grattafiori, W. Xiong, A. Défossez, J. Copet, F. Azhar, H. Touvron, L. Martin, N. Usunier, T. Scialom, and G. Synnaeve, "Code llama: Open foundation models for code," 2023.

[33] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," *Information processing & management*, vol. 24, no. 5, pp. 513–523, 1988.

[34] J. Saxe and K. Berlin, "Deep neural network based malware detection using two dimensional binary program features," *2015 10th International Conference on Malicious and Unwanted Software (MALWARE)*, pp. 11–20, 2015.

[35] S. Schiavoni, F. Maggi, L. Cavallaro, and S. Zanero, "Phoenix: Dga-based botnet tracking and intelligence," in *International Conference*

on detection of intrusions and malware, and vulnerability assessment. Springer, 2014, pp. 192–211.

[36] M. G. Schultz, E. Eskin, E. Zadok, and S. J. Stolfo, "Data mining methods for detection of new malicious executables," in *Proceedings 2001 IEEE Symposium on Security and Privacy. S&P 2001*. IEEE, 2001, pp. 38–49.

[37] M. U. Siddiqui, Y. Wang, Y. H. Lee, and H. Lee, "Novel feature extraction, selection and fusion for effective malware family classification," in *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy*. ACM, 2016, pp. 69–78.

[38] A. Silva, S. Fang, and M. Monperrus, "Repairllama: Efficient representations and fine-tuned adapters for program repair," *arXiv preprint arXiv:2312.15698*, 2023.

[39] Y. Sun, D. Wu, Y. Xue, H. Liu, H. Wang, Z. Xu, X. Xie, and Y. Liu, "Gptscan: Detecting logic vulnerabilities in smart contracts by combining gpt with program analysis," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.

[40] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.

[41] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom, "Llama 2: Open foundation and fine-tuned chat models," no. arXiv:2307.09288, Jul. 2023, arXiv:2307.09288 [cs]. [Online]. Available: http://arxiv.org/abs/2307.09288

[42] Q. Wang, W. U. Hassan, D. Li, K. Jee, X. Yu, K. Zou, J. Rhee, Z. Chen, W. Cheng, C. A. Gunter *et al.*, "You are what you do: Hunting stealthy malware via data provenance analysis." in *NDSS*, 2020.

[43] M. Weyssow, X. Zhou, K. Kim, D. Lo, and H. Sahraoui, "Exploring parameter-efficient fine-tuning techniques for code generation with large language models," *ACM Transactions on Software Engineering and Methodology*, 2023.

[44] B. Wilhelm *et al.*, "A method for summarizing and classifying evasive malware," in *Proceedings of the 26th International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*, 2023.

[45] W. K. Wong, D. Wu, H. Wang, Z. Li, Z. Liu, S. Wang, Q. Tang, S. Nie, and S. Wu, "DecLLM: LLM-Augmented Recompilable Decompilation for Enabling Programmatic Use of Decompiled Code," *Proceedings of the ACM on Software Engineering*, vol. 2, no. ISSTA, pp. 1841–1864, 2025.

[46] D. Wu, D. Gao, R. H. Deng, and C. R. KC, "When program analysis meets bytecode search: Targeted and efficient inter-procedural analysis of modern android apps in BackDroid," in *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2021, pp. 543–554.

[47] C. S. Xia, M. Paltenghi, J. L. Tian, M. Pradel, and L. Zhang, "Fuzz4all: Universal fuzzing with large language models," no. arXiv:2308.04748, Jan. 2024, arXiv:2308.04748 [cs]. [Online]. Available: http://arxiv.org/abs/2308.04748

[48] L. K. Yan and H. Yin, "DroidScope: Seamlessly reconstructing the OS and dalvik semantic views for dynamic android malware analysis," in *21st USENIX security symposium (USENIX security 12)*, 2012, pp. 569–584.

[49] L. Yang, W. Guo, Q. Hao, A. Ciptadi, A. Ahmadzadeh, X. Xing, and G. Wang, "{CADE}: Detecting and explaining concept drift samples for security applications," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 2327–2344.

[50] H. Yin, D. Song, M. Egele, C. Kruegel, and E. Kirda, "Panorama: capturing system-wide information flow for malware detection and analysis," in *Proceedings of the 14th ACM conference on Computer and communications security*, 2007, pp. 116–127.

[51] K. Zhang, Z. Li, D. Wu, S. Wang, and X. Xia, "Low-cost and comprehensive non-textual input fuzzing with llm-synthesized input generators," in *USENIX Security Symposium*, 2025.

[52] L. Zhang, K. Li, K. Sun, D. Wu, Y. Liu, H. Tian, and Y. Liu, "ACFix: Guiding LLMs with Mined Common RBAC Practices for Context-Aware Repair of Access Control Vulnerabilities in Smart Contracts," *IEEE Transactions on Software Engineering*, 2025.

[53] J. Zhao, D. Yang, L. Zhang, X. Lian, and Z. Yang, "Enhancing llm-based automated program repair with design rationales," *arXiv e-prints*, pp. arXiv–2408, 2024.

[54] J. Zhong, D. Wu, Y. Liu, M. Xie, Y. Liu, Y. Li, and N. Liu, "DeFiScope: Detecting Various DeFi Price Manipulations with LLM Reasoning," *arXiv preprint arXiv:2502.11521*, 2025.