

# The Midas Touch: Triggering the Capability of LLMs for RM-API Misuse Detection

Yi Yang<sup>1,2</sup>

Jinghua Liu<sup>1,2</sup>

Kai Chen<sup>1,2,\*</sup>

Miaoqian Lin<sup>1,2</sup>

<sup>1</sup>Institute of Information Engineering, Chinese Academy of Sciences, China

<sup>2</sup>School of Cyber Security, University of Chinese Academy of Sciences, China

{yangyi, liujinghua, chenka, linmiaoqian}@ie.ac.cn

**Abstract**—As the basis of software resource management (RM), strictly following the RM-API constraints guarantees secure resource management and software. To enhance the RM-API application, researchers find it effective in detecting RM-API misuse on open-source software according to RM-API constraints retrieved from documentation and code. However, the current pattern-matching constraint retrieval methods have limitations: the documentation-based methods leave many API constraints irregularly distributed or involving neutral sentiment undiscovered; the code-based methods result in many false bugs due to incorrect API usage since not all high-frequency usages are correct. Therefore, people propose to utilize Large Language Models (LLMs) for RM-API constraint retrieval with their potential on text analysis and generation. However, directly using LLMs has limitations due to the hallucinations. The LLMs fabricate answers without expertise leaving many RM APIs undiscovered and generating incorrect answers even with evidence introducing incorrect RM-API constraints and false bugs.

In this paper, we propose an LLM-empowered RM-API misuse detection solution, *ChatDetector*, which fully automates LLMs for documentation understanding which helps RM-API constraints retrieval and RM-API misuse detection. To correctly retrieve the RM-API constraints, *ChatDetector* is inspired by the ReAct framework which is optimized based on Chain-of-Thought (CoT) to decompose the complex task into allocation APIs identification, RM-object (allocated/released by RM APIs) extraction and RM-APIs pairing (RM APIs usually exist in pairs). It first verifies the semantics of allocation APIs based on the retrieved RM sentences from API documentation through LLMs. Inspired by the LLMs' performance on various prompting methods, *ChatDetector* adopts a two-dimensional prompting approach for cross-validation. At the same time, an inconsistency-checking approach between the LLMs' output and the reasoning process is adopted for the allocation APIs confirmation with an off-the-shelf Natural Language Processing (NLP) tool. To accurately pair the RM-APIs, *ChatDetector* decomposes the task again and identifies the RM-object type first, with which it can then accurately pair the releasing APIs and further construct the RM-API constraints for misuse detection. With the diminished hallucinations, *ChatDetector* identifies 165 pairs of RM-APIs with a precision of 98.21% compared with the state-of-the-art API

detectors. By employing a static detector CodeQL, we ethically report 115 security bugs on the applications integrating on six popular libraries to the developers, which may result in severe issues, such as Denial-of-Services (DoS) and memory corruption. Compared with the end-to-end benchmark method, the result shows that *ChatDetector* can retrieve at least 47% more RM sentences and 80.85% more RM-API constraints. Since no work exists specified in utilizing LLMs for RM-API misuse detection to our best knowledge, the inspiring results show that LLMs can assist in generating more constraints beyond expertise and can be used for bug detection. It also indicates that future research could transfer from overcoming the bottlenecks of traditional NLP tools to creatively utilizing LLMs for security research.

## I. INTRODUCTION

Creating applications requires utilizing Application Programming Interfaces (APIs), and developers are expected to refer to library documentation to ensure their appropriate implementation. Safely utilizing resource-management (RM) APIs, usually employed in pairs (for instance, allocation and releasing), demands strictly adherence to explicitly outlined constraints. Breaching these constraints could lead to security vulnerabilities, including but not limited to memory leaks, use-after-free issues, and double-free security bugs.

Traditional methods use common-used keywords to identify RM-APIs [1]. Recently many other studies have proposed to achieve API misuse detection by analyzing API documentation [2], [3] within which the API constraints are described in strong sentiment or by detecting deviations from frequently used API usages which could be extracted from code analysis [4], [5], [6]. However, all these methods are limited by expertise. Directly using keywords will result in large number of false negatives, the statistics say it will introduce at least 76.28% RM-APIs missing from identification. For those involving documentation analysis, it is hard to retrieve the API constraints beyond the documentation; For those focusing on code analysis, the incorrect usage will influence the whole detection procedure introducing false bugs. Traditional program analysis also has difficulties identifying if the current allocation function needs extra releasing operations within multi-layered nesting. The statistics of the GroundTruth in Section V shows 72.7% of functions have multi-layered nesting and merely 27.3% of functions directly call `malloc/free` functions. These results show that traditional program analysis will introduce false negatives in RM-API pairing due to complex

\* Corresponding Author

code structures; it also introduces false positives when the allocated/released object has already been released within the function and no extra releasing operations are required. Therefore, with the potential of Large Language Models (LLMs) on text analysis and generation, researchers tend to explore the ability of LLMs on bug detection.

**Challenges in exploiting LLMs for RM-API misuse detection.** We attempt to use one of the state-of-the-art LLMs to retrieve information from API official documentation, identify resource allocation functions, and pair them with the corresponding releasing functions. However, directly using LLMs still encounters the following challenges.

**C1: LLMs fabricate answers without expertise.** We observed that when directly prompting LLMs to provide corresponding releasing functions based on the existing allocation functions, LLMs tend to fabricate answers according to the naming format. For example, `evwatch_check_new` and `evwatch_free` form a pair of RM-APIs, but LLMs provide the resource-releasing function `evwatch_check_free`, which follows a naming pattern similar to `evwatch_check_new`, and the operation part includes a typical antonym pair, i.e., *new* and *free*. However, this function does not exist, and LLMs exhibit obvious fabrication probably due to the absence of information. Therefore, supplementing valid information and designing the solution procedure are crucial in addressing the fabrication issues.

**C2: LLMs introduce incorrect answers with evidence.** Another challenge arises from LLMs producing incorrect answers even with grounded evidence. As a black-box model, when identifying the current function’s semantics based on specific sentences with LLMs, we cannot guide the model to output the expected results. For example, after prompting LLMs with the RM sentence “*The zip\_discard() function closes the archive and frees the memory allocated for it*” of `zip_discard`, and prompting LLMs to determine if the current function has allocation semantics. Even though the RM sentence contains the keyword *free* indicating an obvious resource-releasing semantics, LLMs still provide the incorrect answer *YES*. This phenomenon is quite common and can impact the final identification of allocation functions, potentially leading to false bugs. Therefore, solving the encountered issues in LLMs and generating accurate answers is also important.

**ChatDetector.** To address the challenges mentioned above, we propose an LLM-empowered RM-API misuse detection solution, *ChatDetector*. The challenges mentioned above happen in most of the state-of-the-art LLMs (see Section VI) and we perform *ChatDetector* based on ChatGPT considering the overall performance. We adopt the inspiration from ReACT framework [7] and compose *ChatDetector* including three components: 1) utilizing ChatGPT to identify resource allocation functions from API documentation; 2) employing ChatGPT to provide corresponding resource-releasing functions for the identified resource allocation functions; 3) using CodeQL together with the identified RM-API pairs to detect misuses in open-source software.

First, *ChatDetector* decomposes the complex task into two

steps: identifying RM (Resource-management) sentences from API official documentation that describe the current function’s functionality and are related to resource-allocation operations; and then, determining which function has resource allocation semantics based on the RM sentences. Specifically, the recognized RM sentences come in two types: 1) RM sentences that contain the function name and the allocated object (RM object); and 2) RM sentences that only include the allocated object. According to the various prompting methods and information, we design a two-dimensional prompting method to get the reasoning answers for allocation APIs identification. After that, *ChatDetector* employs cross-validation for the LLMs’ output to identify the potential allocation APIs. Through result inspection, we discover that the current method still prompts LLMs to produce evident errors in recognizing function semantics based on RM sentences. For example, when RM sentences contain obvious keywords indicating resource-releasing semantics like *free*, ChatGPT still erroneously identifies the current function’s semantics as resource allocation. This obvious inconsistency between the input and output results in a significant number of false positives. To address this issue, we propose an inconsistency check between the LLMs’ output and the reasoning process. Specifically, when the current function is identified as having allocation semantics, *ChatDetector* uses Semantic Role Labeling [8] (an off-the-shelf NLP tool) to extract the operation keywords corresponding to the current function during the reasoning process. Then *ChatDetector* will classify the operation into *allocate*, *free* or *neutral* to further validate that the current answer is generated based on the original text rather than fabricated. Through cross-validation of answers provided after the three prompts and inconsistency checks on the reasoning process, *ChatDetector* can recognize resource allocation functions and their corresponding descriptions of RM objects.

Then, *ChatDetector* attempts to prompt LLMs to provide the corresponding resource-releasing functions based on the identified resource-allocation functions. Due to the same issues encountered above, LLMs also tend to fabricate answers based on the function’s naming format due to incomplete information and complex tasks. Therefore, based on the fact that both resource-management functions act on the same RM object, *ChatDetector* decomposes the pairing task, which is inspired by the ReAct framework, into RM-object type extraction and releasing API pairing. Specifically, *ChatDetector* locates the corresponding RM-object type in the current function declaration based on the RM object description. Due to the poor performance of LLMs in identifying object types based on unstructured and lengthy function declarations, *ChatDetector* pre-processes the function declarations into formatted data, distinguishing the return type and name as well as the types and names of each parameter, which enables LLMs to better identify RM-object types. Based on the identified RM object types, *ChatDetector* can generate accurate resource-releasing functions with diminished “hallucinations” in LLMs. This illustrates that the current approach, combining external knowledge and step-by-step reasoning, can enhance the accu-

racy of LLMs in identifying RM-API pairs.

Finally, we evaluate *ChatDetector* using the following research questions:

**RQ1:** Can *ChatDetector* be utilized for RM-API misuse detection? **RQ2:** Can *ChatDetector* discover more RM-API pairs than the previous work? **RQ3:** Is it necessary to perform reasoning with LLMs for RM-API misuse detection?

On the dataset including six libraries (FFmpeg [9], Libevent [10], Libexpat [11], Libpcap [12], Libzip [13] and OpenLdap [14]), *ChatDetector* can identify 165 RM-API pairs, with relatively high precision of 92.81% in allocation APIs identification and 98.21% in RM-API pairing, outperforming the benchmark work by identifying 80.85% more RM-API pairs. *ChatDetector* discovers 115 security issues which were ethically reported to the developers (in Section IV-A).

**Contributions.** Our contributions are listed as follows:

- **Novel research direction.** To the best of our knowledge, *ChatDetector* is the first to propose using LLMs to understand and extract API usage patterns from official API documentation, enabling misuse detection in real software. This work opens up new research directions for bug detection aided by large models, extending beyond the bottlenecks of traditional NLP tools to explore how to effectively leverage LLMs for security research.
- **New approach.** We introduce an automated tool that leverages LLMs to extract API usage patterns from official documentation. We are inspired by the ReAct framework to decompose the RM-API constraint retrieval into two parts with additional expertise: 1) allocation APIs identification which adopts a two-dimensional prompting method for cross-validation and an inconsistency-checking approach between the answer and the reasoning process; 2) releasing APIs pairing which adopts pre-processed function declaration for diminishing hallucinations. Compared with the state-of-the-art API misuse detectors, *ChatDetector* is able to discover at least 80.85% more RM-API pairs beyond expert knowledge and achieved an accuracy of 98.21%. When applied to six popular libraries and the applications that use them, *ChatDetector* discovered 115 security issues and were ethically reported to the developers.
- **Insightful Findings.** Through the analysis of the detection results of *ChatDetector*, we found that functions with lower usage frequency correspond to relatively incomplete documentation descriptions. Additionally, we observed that among the resource-allocation functions identified by *ChatDetector*, only 52.17% of the functions had RM sentences with commonly used keywords. Although the inspiring results show that LLMs has assisted in discovering more constraints beyond expertise, the small number of false positives introduced by this black-box model is currently challenging to explain which needs to be addressed in future research.

## II. BACKGROUND AND PRELIMINARY STUDY

### A. Background

**Large Language Models.** Large Language Models (LLMs), trained on vast corpora, exhibit proficiency across diverse

natural language processing tasks, including common sense question-answering and semantic extraction [15]. The emergence of models like ChatGPT [16] has brought heightened attention to both the security issues inherent in LLMs and their potential applications in addressing security challenges. Among these, ChatGPT stands out as a state-of-the-art model designed for general-purpose tasks, initially based on the GPT-3.5 architecture and even more effective GPT-4 [17]. Taking both cost and performance into consideration, we opted to employ GPT-3.5 for implementation.

**Reasoning with Large Language Models.** Initially, recognizing the limitations of LLMs in tackling complex problems, researchers introduced the concept of the *Chain-of-Thought* (CoT) approach [18]. This methodology guides LLMs through a step-by-step thought process, empowering them to effectively address complicated problems and generate precise answers. However, conventional CoT methods have exhibited considerably lower effectiveness in resolving textual reasoning tasks in comparison to math tasks. Consequently, researchers have introduced a novel framework known as ReAct framework [7]. By combining CoT with external knowledge, the “*Think-Act-Observe*” framework is employed for task decomposition. Inspired by these, our approach breaks down the complex task of recognizing RM-API constraints by supplying additional documentation content to the LLMs. This augmentation enables efficient identification and precise matching of RM-APIs.

**CodeQL.** CodeQL [19] is a static analysis tool for detecting security issues with QL queries written by developers or on users’ own. The analysis procedure consists of three steps: creating a CodeQL database, running QL queries against the database, and interpreting the analysis results. CodeQL generates a database by extracting relational representations of source files, monitoring build processes for compiled languages, and resolving dependencies for interpreted ones. Each supported language has its extractor to maintain accuracy, and databases are created one language at a time before importing data into a CodeQL database directory. Once a CodeQL database is created, developers execute QL language queries against it. In the final step, results from query execution are interpreted to highlight potential issues in the source code, based on metadata properties within the queries.

### B. Preliminary Study

In this section, we aim to tackle two tasks: allocation-API identification and RM-API pairing, directly leveraging large language models (LLMs). We explicitly outline the challenges and specific issues encountered when employing this approach.

① **LLMs fabricate answers without expertise.** We initially prompt LLMs for identifying RM APIs, as depicted in Figure 1. LLMs responds that it finds it difficult to provide an answer due to insufficient information. Upon refining the prompt, LLMs indicates that the current API `zip_file_set_comment` is not involved in allocation. However, upon inspecting the source code, it is revealed that `zip_file_set_comment` indirectly invokes `malloc`,

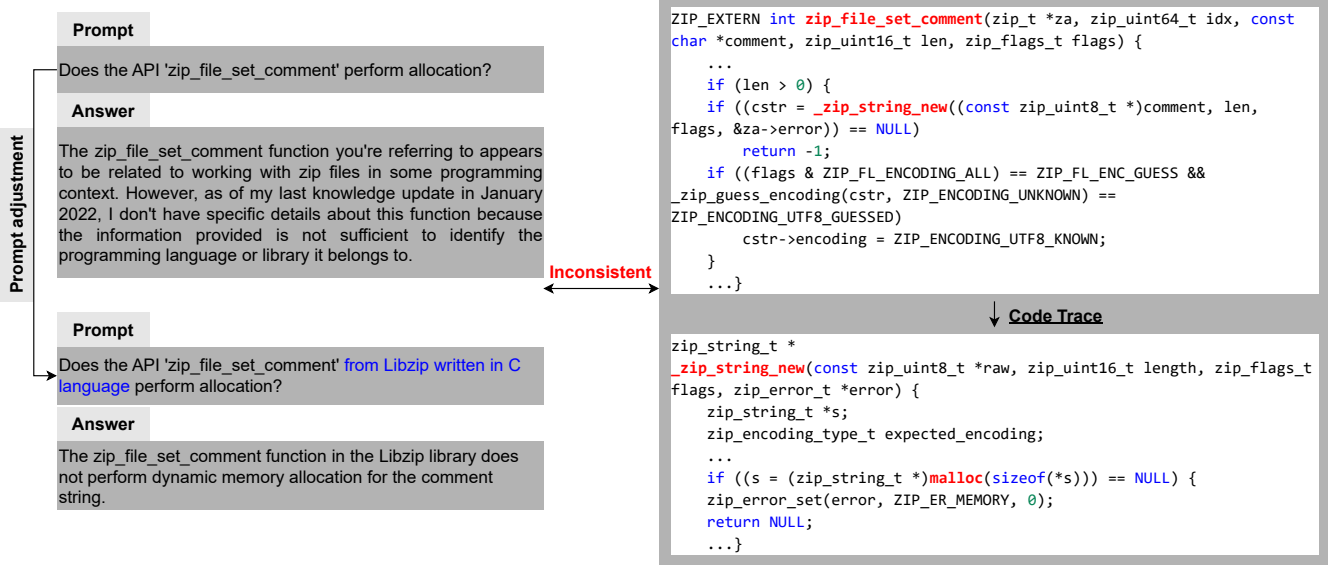


Fig. 1: The limitation for direct-use of LLMs for RM-API Identification

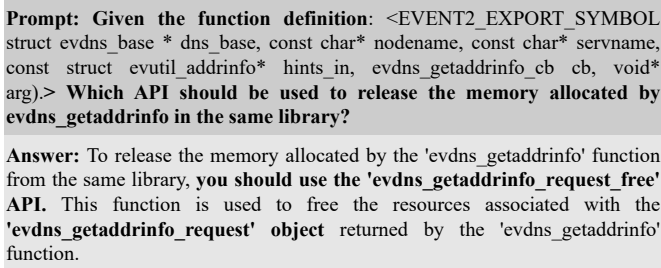


Fig. 2: The limitation for direct-use of LLMs for RM-API pairing

indicating its allocation semantics, contrary to the assertion made by LLMs. The same phenomenon also exists in RM-API Pairing. We directly utilize LLMs for releasing API retrieval as illustrated in Figure 2. When providing the function declaration of `evdns_getaddrinfo` to LLMs and asking for the corresponding resource-releasing API, the LLMs outputs the result `evdns_getaddrinfo_request_free`. As shown in the figure, LLMs not only gives the *function description* of this API but also provides the object type that needs to be released (RM Object Type). This seemingly reasonable answer and explanation are, in fact, incorrect. The truth is that the function `evdns_getaddrinfo_request` does not exist and was fabricated by LLMs. Such a phenomenon could also be one of the impacts introduced by LLMs’s hallucinations. Due to the lack of expertise, LLMs fabricates answers based on prior experience or stereotypes. For instance, LLMs assumes that a resource-releasing API should contain the keyword *free*, and the corresponding API should contain partially identical keywords such as `evdns_getaddrinfo`. Therefore, we propose to perform reasoning with LLMs and provide them with external expertise to improve the accuracy of downstream tasks.

⊗ **LLMs generate incorrect answers with evidence.** Another challenge occurs in the reasoning process, for example,

when providing the descriptions of `zip_discard` “*The zip\_discard() function closes the archive and frees the memory allocated for it*” to LLMs, the result of the prompt *Does this API perform allocation?* is “YES” which is contrary to the fact. One possible reason for these cases is that hallucinations exist in LLMs [20]. According to the recent research [21], the reasons for this phenomenon are typically a lack of context, leading to an inability to understand the questions or a lack of expertise resulting in inaccurate answers. Therefore, we propose an inconsistency-checking approach between the LLMs’ output and the reasoning process to achieve better performance in RM-API Identification.

### III. DESIGN AND IMPLEMENTATION

#### A. Overview

In this section, we provide a detailed explanation of our approach, *ChatDetector*, which can analyse official documentation with LLMs to extract RM-API constraints and achieve RM-API misuse detection. We also provide a description of the overall framework and then illustrate each component with detailed examples.

**Framework.** This work aims to explore the capabilities of LLMs in RM-API constraints retrieval and RM-API misuse detection. As illustrated in Figure 3, *ChatDetector* consists of three main components: utilizing LLMs to analyze documentation for resource-allocation APIs identification (Figure 3 (a-c)), recognizing corresponding resource-releasing APIs with LLMs (Figure 3 (d-f)), and performing RM-API misuse detection using CodeQL based on the identified RM-API pairs (Figure 3 (g)). *ChatDetector* adopts ChatGPT (GPT-3.5-turbo model [22]), one of the state-of-the-art LLMs due to its superior performance and cost-effectiveness (detailed discussion are shown in Section VI), for the whole implementation. The whole process of LLMs utilization is fully automated and

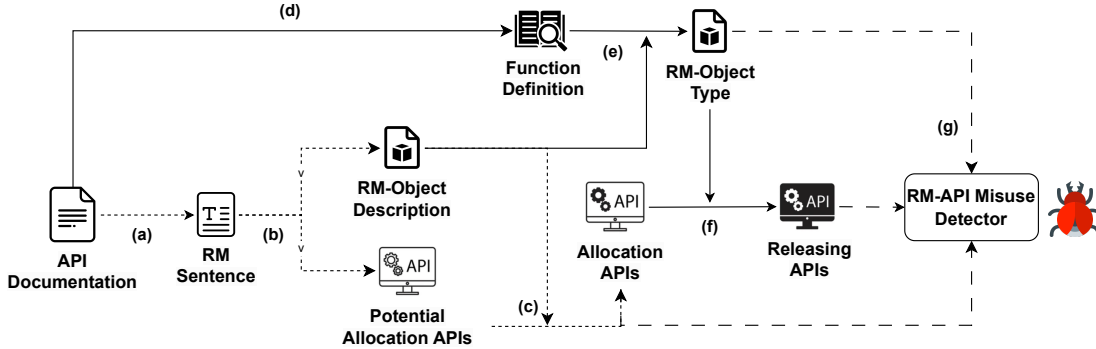


Fig. 3: Architecture of *ChatDetector*

TABLE I: Examples of RM Sentences

	RM Sentences	RM API	RM object Description
Type-I	“The functions <code>zip_source_filep()</code> and <code>zip_source_filep_create()</code> create a <b>zip source</b> from a file stream.”	<code>zip_source_filep()</code> and <code>zip_source_filep_create()</code>	a zip source
Type-II	“a <b>newly allocated struct event</b> that must later be freed with <code>event_free()</code> or NULL if an error occurred.”	-	struct event

human assistance is only needed in detection code generation when using CodeQL for misuse detection, which is a one-time effort. More specifically, inspired by ReAct framework, *ChatDetector* decomposes the allocation-API identification into two parts. First, it identifies the RM sentences from the official documentation describing the current API’s functionality and possessing resource-allocation semantics. Then, adopting a two-dimensional prompting method for cross-validation, *ChatDetector* prompts to provide potential resource-allocation APIs and descriptions of the allocated objects (RM objects). Through the inconsistency-checking between the LLMs’ output and the reasoning process, APIs with allocation semantics are obtained. Next, by utilizing the pre-processed function declarations, *ChatDetector* can extract the corresponding RM object types and output the corresponding resource-releasing APIs with diminished hallucinations. Finally, the identified RM-API pairs are written into QL queries, and RM-API misuse detection is performed using CodeQL [19] on open-source software.

### B. RM-API Identification

**API Semantics Confirmation.** We provide a detailed definition for the RM-APIs. Specifically, we consider the APIs that call the `malloc` function to have allocation semantics; and the APIs that call the `free` function to have releasing semantics. We manually check each API’s source code and documentation according to the procedure proposed in Section V. To further confirm the allocation semantics of the API, we obtain the allocated object of the current API through data-flow analysis and control-flow analysis. If the allocated object needs an extra releasing operation outside the API, we assume the current API has allocation semantics. The semantics of the specific API is confirmed by the latest official documentation and source code. Besides, the semantics of a function is fixed from the outset and in most cases does not change due to iteration.

In other words, when a function is identified to have allocation semantics, its functionality will not change due to the iteration of function versions, that is, it will not switch from having allocation semantics to having releasing semantics.

**The Definitions.** Through analysis of a large amount of API documentation, we observe that APIs with **RM semantics** (such as, allocate/release) usually have explicit or implicit descriptions within the documentation. We term the sentences that have the semantics of such APIs as **RM sentences**. More specifically, if an API calls `malloc/free` in the source code and its description includes the keywords that are semantically similar to `malloc/free`, we treat this sentence as RM sentences, the keywords that represent `malloc` semantics fall in “create”, “allocate”, “construct” and “initialize” etc. We then construct the GroundTruth of 244 sentences through manual inspection from three libraries including FFmpeg, Libdbus and Libpcap. The statistics show that 18.4% of RM sentences contain function names and 81.6% of them contain the allocated/released objects. Based on this, we classify the RM sentences into two types, as illustrated in Table I. The first type of RM sentence is extracted from the description of the function `zip_source_filep`. This sentence explicitly states that the function containing memory allocation functionality and the memory allocated is stored in “a zip source”, which we term as the **RM object description**. The second type of RM sentence is extracted from the return value description of the function `event_new`. This sentence implies the memory allocation semantics of the function by stating that the return value `struct event`, which can be referred to as **RM object**, represents the allocated memory and requires releasing using another function. Although extracting the RM-API name directly from this sentence is challenging, we can infer the RM object described in this sentence by understanding its semantic context.

From the preliminary study in Section II-B, directly using

Given the description of the API '{api}' is: {desc}. Output the following answers in JSON format.

1. Is there a sentence describing the semantics of the API '{api}' related to allocation or releasing in the text? If YES, please output the first statement that meets the criteria; if NO, please output NULL.
2. Output the evidence sentence grounded for getting the result in '1'.
3. Output the reasoning process for getting the answers in '1'.

Fig. 4: The prompt for extracting RM Sentences

Given the RM sentence of the API '{api}' is : {rm\_sents}  
Output the following answers in JSON format.

1. Which API performs allocation? Output the API name only.
2. Present the evidence for getting this API.
3. Present your reasoning process for getting this API.

Fig. 5: The prompt (a) for identifying RM-APIs

LLMs for RM-API identification results in many incorrect answers. We also explored adding more information to identify API's semantics with source code or documentation (shown in Section IV-C), but neither achieved satisfactory results in the current task. Therefore, inspired by the ReAct framework [7] optimizing on *Chain-of-Thought* [18], we propose to decompose the complex task (RM-API constraint retrieval) into RM sentence extraction and allocation APIs identification, together with additional information retrieved from API documentation.

**RM Sentence Extraction.** According to the definitions above, RM sentence describes the functionality of the RM APIs and *ChatDetector* uses the prompt in Figure 4 to extract them from API documentation descriptions. Note that all the content such as the *api* and *desc* representing description, highlighted in brackets are automatically crawled from official documentation. The rest context in the prompt is the template which is a one-time effort. The procedure is illustrated as follows. First, *ChatDetector* prompts the LLMs to provide the first sentence in the documentation describing the current API's functionality and check if it possesses RM semantics. When multiple API descriptions are present on the same documentation page, the sentence outlining a particular API's functionality typically resides close to that API. In cases where documentation exclusively details a single API and contains multiple sentences with RM semantics, to prevent redundancy, we exclusively select the first sentence that delineates the API's functionality and exhibits RM semantics as the RM sentence. Secondly, owing to the potential issue of *hallucination* in the output of LLMs [23], it is imperative to verify whether the presently identified RM sentence originates from the documentation itself and is not fabricated. Furthermore, confirming the accuracy of this response is essential. Consequently, *ChatDetector* prompts LLMs to provide the reasoning process and the source-based evidence of RM sentence identification to support the answer.

**Allocation APIs Identification.** To perform allocation APIs

Given the RM sentence of the API '{api}' is : {rm\_sents}  
Output the following answers in JSON format.

1. Does the API performs allocation? Answer YES/NO.
2. Present the evidence for getting the answer in '1'.
3. Present your reasoning process for getting the answer in '1'.

Fig. 6: The prompt(b) for identifying RM-APIs

Given the RM sentence of the API '{api}' is : {rm\_sents}  
Output the following answers in JSON format.

1. Output the allocated object according to the RM sentence.

Fig. 7: The prompt for identifying RM-objects

identification, we tried to directly apply the extracted RM sentences to the prompt. However, due to the inherent uncertainty introduced by LLMs, accurately validating the truthfulness and accuracy of current answers using a single prompt poses challenges. We observed the RM sentences can be classified into two categories: those containing both RM-API names and RM objects, and those containing RM objects only. For sentences containing RM objects only, RM APIs identification can be achieved by assessing the presence of RM objects and verifying the accuracy of LLMs's response. On the other hand, for those sentences containing both RM-API names and RM objects, RM APIs identification necessitates additional validation for authenticity and accuracy, building upon the previous step. And also, the previous studies show that LLMs are effective in three types of prompting methods, which are *in-context prompting*: extracting answers from the given context and the answers can be extracted from the input context; *open prompting*: generating answers based on the information from the given context and the answers cannot be extracted from the input context; and *indirect prompting*: generating answers based on the indirect questions, both the prompts and the answers pose indirect relations to the final answers. Therefore, *ChatDetector* follows the above three types of prompting and adopts the observation to design the prompts in Figure 5 and Figure 6 to prompt LLMs for allocation APIs identification based on RM sentences, and utilizes the prompt in Figure 7 to obtain RM-objects description.

Specifically, the prompt in Figure 5 is designed to inquire about the API with resource allocation functionality based on the reasoning derived from the current RM sentence. Additionally, the prompt in Figure 6 is designed to ascertain whether the current API possesses resource allocation semantics. And the prompt in Figure 7 is intended to acquire the RM-object. To further confirm the allocation APIs, we design rules (details are shown in Figure 8) to cross-validate the three answers. The necessary condition is to ensure the extracted RM objects exist in the documentation; if the response to the prompt in Figure 5 is the `current API` or the response to the prompt in Figure 6 is YES, then *ChatDetector* collects the potential allocations APIs through the cross-validation.

However, the current method leads to many false positives, primarily due to ChatGPT being a black-box model, making

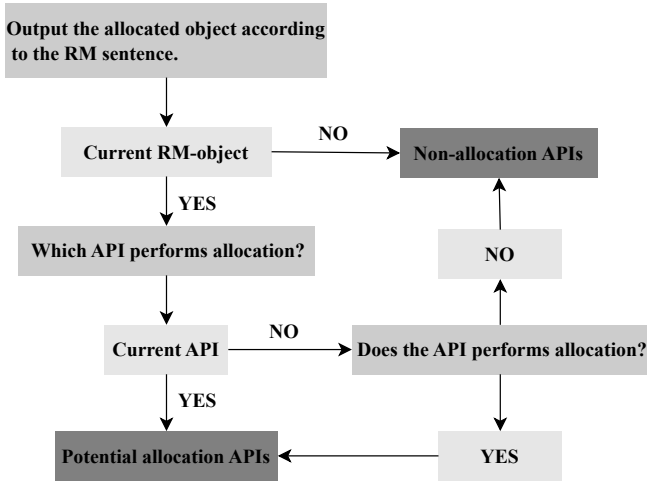


Fig. 8: The rules for Cross-Validation

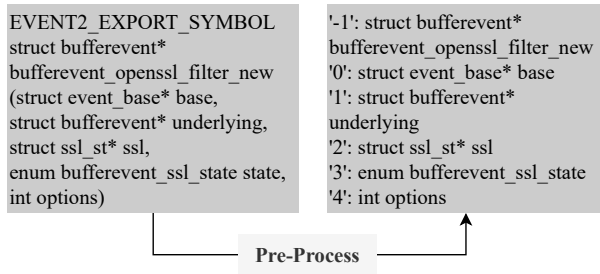


Fig. 9: The pre-process procedure of Function Definition

it challenging to be guided. Besides, we observe typical self-inconsistencies between the LLMs’s output and the reasoning processes. For example, LLMs indicate that `zip_close` has memory allocation semantics. But the reasoning process for `zip_close` is as follows: “The sentence mentions that the `zip_close` function **fre**s the archive, indicating that its purpose is to **deallocate** resources rather than allocate them. Therefore, it can be inferred that the function `zip_close` **does not perform allocation**.” The reasoning process confirms the identified RM sentence is correct, but LLMs has limitations in understanding this sentence. Even when the RM sentence clearly states that the current function’s operation is memory deallocation, the model still interprets the function’s semantics as the opposite. Therefore, we utilize an off-the-shelf NLP tool, Semantic Role Labeling [24], to extract the action verbs corresponding to the current function within the reasoning process, such as *set*, *free* and others. By comparing the semantic similarity between these action verbs and *release*, *allocate* or *Neutral*, *ChatDetector* can further determine the semantic tendency of the current function and acquire the resource-allocation APIs.

### C. RM-API Pairing

According to the preliminary study in Section II-B, we attempted to directly apply the obtained allocation APIs along with the function definition for corresponding releasing APIs identification by using the prompt in Figure 12. However, due to the inherent *hallucinations* in LLMs, a great number of

Given the allocated object description of '{api}' is '{obj\_desc}' and the function definition is '{func\_def}'. Output the following answer in a JSON format.

1. According to the function definition, output the object type
2. which has been best described by the object description.
3. Along with the number quoted

Fig. 10: The prompt for RM-object Identification

Given the allocated object type of the API '{malloc\_api}' is '{object\_type}' according to the function definition '{func\_def}'.

1. Which API that has the same object type should be used to release the allocated memory of '{malloc\_api}'?

Fig. 11: The prompt for Releasing API Identification

fabricated answers were generated. LLMs might fabricate the answers due to the inability to solve complex tasks and insufficient information for simple task completion. The common knowledge says that the RM-API pairs share the same RM-object type. That is, the released object of the releasing API is typically the allocated object of the allocation API, the naming of these two APIs may differ, but the type of RM object remains consistent. Based on this knowledge, *ChatDetector* adopts the inspiration from ReAct framework to decompose the releasing APIs identification task into RM-object Type Extraction and releasing API Identification.

**RM-object Type Extraction.** To achieve RM-object Type Extraction, we first attempted to prompt LLMs with RM object descriptions retrieved from Section III-B and function declarations which were automatically retrieved from API documentation. However, this approach did not effectively enable LLMs to identify the RM object’s type. For example, when trying the prompt in Figure 12 for releasing API identification, the output of LLMs tended to confuse the function names and types, which may result from insufficient information in the training set of the model. As a solution, *ChatDetector* pre-processes the function declarations by separating the types and names of parameters/return values, the details are illustrated in Figure 9. Specifically, for `bufferevent_openssl_filter_new` with its complete declaration, *ChatDetector* pre-processes it by splitting the types and names of parameters and return values along with numeric markers. The pre-processed declaration is shown in the figure. Further, *ChatDetector* uses the prompt in Figure 10 to extract the corresponding RM object type. Experiments have shown that breaking down lengthy function declarations into parts can enhance ChatGPT’s accuracy in matching RM object types based on the function declaration.

**Releasing API Pairing.** Based on the extracted RM object type, *ChatDetector* further acquires the corresponding resource-releasing APIs by providing LLMs with the identified allocation API and the RM-object type with the utilization of the prompt in Figure 11. In this step, we did not design a complex solution. The reason is that the *hallucination* issues observed in the initial design of LLMs on the resource-allocation APIs Identification gradually disappeared as we



Given the function definition of '{api}' is: '{func\_def}'  
Which API should be used to release the memory allocated by '{api}'  
in the same library?

Fig. 12: The prompt for initial releasing API Identification

continuously provided valid expert knowledge and complex task decomposition solutions. Experiments in Section IV have shown that *ChatDetector* can accurately identify the correct RM-API pairs. This result indicates that *hallucination* can be further resolved through continuous interaction with expertise and the decomposition inspired from *Chain-of-Thought*.

#### D. RM-API misuse detection

Based on the previously identified RM-API pairs, we have obtained resource-allocation APIs, resource-releasing APIs, and the shared RM object types. In this section, we will provide a detailed explanation of how to perform RM-API misuse detection using CodeQL [19]. For each RM-API pair, *ChatDetector* applies CodeQL to detect three types of security bugs: memory leaks, use-after-free, and double-free. Applying CodeQL to security bug detection includes three steps (Section II-A). First, we create the database for each application integrating one of the six libraries (details are shown in Table V in Appendix). Then we generate corresponding QL queries based on the detected RM-API pairs (each pair corresponds to one piece of constraint) to detect RM-API misuse. Note that the queries are manually constructed based on the basic template provided by CodeQL and the construction procedure is a one-time effort which requires no need to modify if applied to detect the same bug type on new software. We then interpret the results, confirm the misuse and ethically report to developers. For memory leak bugs, the tool locates a given allocation API and the allocated object. Based on the data-flow analysis and control-flow analysis, the tool checks every path within the code block to see if there is a corresponding releasing API to free the allocated object. If there are paths that do not contain the corresponding releasing API, then a memory leak occurs. For double-free bugs, the tool checks if one path contains two releasing APIs on the same object and there are no other allocation APIs between them which results in double-free bugs. For use-after-free bugs, the tool checks whether there is an operation to dereference the freed object after the releasing API, if so, it is a use-after-free bug.

Take the memory-leak detection strategy of `evconnlistener_new_bind` as an example and the QL code is listed in Listing 1. CodeQL first calls the function `isSourceFC` to locate the control flow node by the function name (`evconnlistener_new_bind`), indicating the current block contains the function. Then the tool calls the function `isLocalSource` to verify the RM object is a local variable. Finally, CodeQL calls the function `getLeakBlock` which internally calls the function `getSourceExpr` to locate the node corresponding to the specified parameter of the function. The method detects no post path of the function

`evconnlistener_new_bind` contains the corresponding releasing function which indicates there exists a memory leak within the basic block. CodeQL supports modifying the RM-API names and the corresponding indexes of RM objects to automatically generate corresponding detection code which helps in large-scale bug detection.

```

1 predicate isSourceFC(FunctionCall fc)
2 {
3 // malloc function: evconnlistener_new_bind
4 fc.getTarget().hasName("evconnlistener_new_bind")
5 }
6
7 Expr getSourceExpr(FunctionCall fc)
8 {
9 //parameter 1:
10 result = fc.getArgument(0)
11 //return value:
12 //result = fc
13 }
14
15 import cpp
16 from BasicBlock bb, FunctionCall malloc
17 where
18 // locate the malloc function call by allocation
19 // API name
20 isSourceFC(malloc)
21 //Make sure that the malloc function operates on
22 // a local variable
23 and isLocalSource(malloc)
24 //There is a path in the current malloc function
25 // postpath that does not have a releasing
26 // operation
27 and bb = getLeakBlock(malloc)
28 select malloc, malloc.getLocation().toString()

```

Listing 1: QL code for detecting the memory leak of function `evconnlistener_new_bind`

## IV. EVALUATION

### A. RQ1: Effectiveness

In this section, we implement *ChatDetector* on open-source libraries and the applications integrated on them to further evaluate *ChatDetector*'s performance.

**End-to-end Effectiveness.** We implement *ChatDetector* on six popular libraries documentation: FFmpeg [9], Libevent [10], Libexpat [11], Libpcap [12], Libzip [13] and OpenLdap [14] (details are shown in Table V in the Appendix), and the Top 10 applications integrating on them. In total, *ChatDetector* discovered 165 RM-API pairs containing 165 allocations APIs and 61 releasing APIs, for some of the allocation APIs may have multiple releasing APIs to release different parameters and some of them may have the same releasing API for multiple allocation APIs. We followed the ethical bug disclosure policy of software and reported all of the 115 security bugs, 64 of which were confirmed by the developers with a precision rate of 90.2%. And the disclosure procedure has no ethical issues. The details are shown in Table VI in the Appendix. Figure 13 illustrates a false positive case in CodeQL's detection when inspecting the use of `zip_source_buffer`. *ChatDetector* detected the correct RM-API pair: the resource-allocation API `zip_source_buffer` and the resource-releasing API



```

void SetFileData(std::string_view path,
std::vector<uint8_t>&& data) override
{
    _writeBuffers.push_back (std::move(data));
    const auto& writeBuffer =
*_writeBuffers.rbegin();

    auto source = zip_source_buffer(_zip,
writeBuffer.data(), writeBuffer.size(), 0);
    auto index = GetIndexFromPath(path);
    zip_int64_t res = 0;
    if (index.has_value())
    {
        res = zip_file_replace(_zip,
index.value(), source, 0);
    }
    else
    {
        res = zip_file_add(_zip,
path.data(), source, 0);
    }
    if (res == -1)
    {
        zip_source_free(source);
        throw std::runtime_error
(std::string(zip_strerror(_zip)));
    }
}

```

Fig. 13: A false positive introduced by CodeQL

zip\_source\_free. CodeQL’s detection revealed that no subsequent call exists to zip\_source\_free to release the allocated resource after calling zip\_source\_buffer. However, as depicted in the figure, the parameter *res* is generated by the functions zip\_file\_replace or zip\_file\_add, and the documentation explicitly states, “NOTE: zip\_source\_free(3) should not be called on a source after it was used successfully in a zip\_file\_add or zip\_file\_replace call.” Consequently, in the condition “res == -1”, which corresponds to a call failure, zip\_source\_free is invoked to release resource. Due to CodeQL’s inability to handle vulnerability detection when multiple functions collaborate, it resulted in a false memory leak detection.

**Effectiveness of RM-API Identification.** ChatDetector detects 168 resource allocation APIs across the six software libraries with an accuracy of 92.81% and a recall of 92.06%. 14 false negatives generated during the identification process were primarily due to the prior misidentification of RM sentences. This could be attributed to LLMs’ difficulty in locating key sentences in less obvious locations. For example, the documentation description for the function zip\_open\_from\_source is illustrated in Figure 21 (details are shown in Appendix). In this description, the sentences related to the zip\_open\_from\_source function do not appear at the beginning of the paragraph or are not prominently highlighted. This leads LLMs to provide the incorrect RM sentence “The semantics of the API zip\_open\_from\_source related to allocation or releasing is not described in the given description”. However, when we prompt LLMs to determine the function’s semantics based on the sentence “The zip\_open\_from\_source() function opens a zip archive encapsulated by the zip\_source zs using the provided flags. In case of error, the zip\_error ze is filled in”, LLMs is able

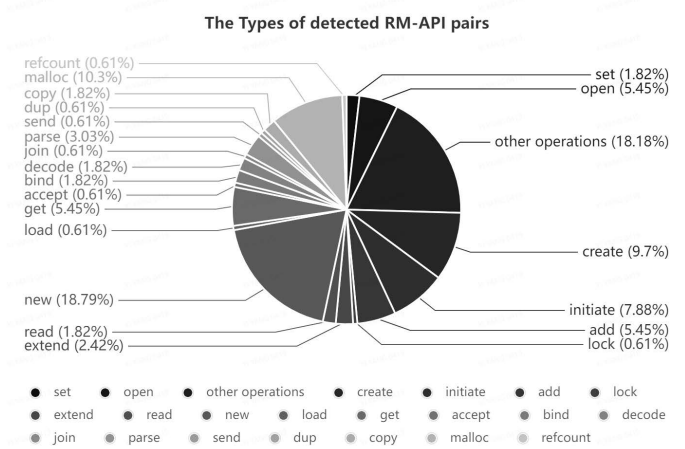


Fig. 14: The types of detected RM-API pairs

to correctly identify that the current function has allocation semantics. Therefore, we attribute the occurrence of these false negatives to the inherent limitations of LLMs’ capabilities.

Simultaneously, ChatDetector also introduced 13 false positives. Through the analysis of these cases, we found that the content in the documentation can also mislead LLMs to make incorrect judgments. For instance, in the documentation description of the function zip\_source\_begin\_write\_cloning (shown in Figure 20 in the Appendix), the phrase “Usually this involves creating temporary files or allocating buffers” explicitly indicates that the current function involves resource allocation operations. However, upon inspecting the source code, it was found that the function did not call any other functions involving malloc. Therefore, the inconsistency between the documentation description and the source code can also lead to false positives in ChatDetector, making these cases of misleading documentation errors.

**Effectiveness of RM-API Pairing.** ChatDetector detects 165 correct RM-API pairs in six software libraries achieving an accuracy of 98.21%. The three false positives encountered during the matching process were primarily due to the lack of common sense in the LLMs and the hallucinations of the model. For example, in the function evhttp\_uri\_join, which allocates memory and stores it in the return value, with the corresponding RM object type char\*. However, even though ChatDetector accurately identified the RM object type, for LLMs, it was difficult to find the corresponding resource release function based on this general type. Therefore, the answer given by ChatDetector was evhttp\_uri\_free, which has a similar naming pattern to the function evhttp\_uri\_join, which however, was fabricated by the model. In the case of the function evwatch\_check\_new, which allocates memory and stores it in struct evwatch\*, LLMs provided a fabricated answer, evwatch\_check\_free, based on the correct RM object type. While this function had a similar naming pattern, it did not actually exist. Although such cases are not common in the results of the current method, these random false positives

TABLE II: Common RM-API pairs and bugs detected from four tools

Library	Our work		Advance		HERO		Goshawk	
	Pair	Bug	Pair	Bug	Pair	Bug	Pair	Bug
Libevent	56	9	4	3	1	0	6	4
Libzip	23	2	2	0	0	0	0	0
Libexpat	6	0	2	0	0	0	0	0
<b>Total</b>	<b>85</b>	11	8	3	1	0	6	4

also result in additional time costs in the detection process.

**The newly discovered RM-API misuses.** Table VI in the Appendix shows that *ChatDetector* discovers 115 security bugs on the applications integrated on six popular libraries, with 165 detected pairs and the details are shown in Figure 14. More specifically, we categorize the RM-APIs based on the verbs extracted from their functionalities. The APIs labelled as “other operations” cannot have their corresponding actions directly extracted from the function names, such as `pcap_list_datalinks`. From the figure, it can be seen that APIs with the “new” operation account for as much as 18.79%, making it the most prevalent. APIs with the “malloc” operation account for 10.3%, ranking second. Other relatively common operations include “create”, “initiate”, “open”, “get” and “add” have equal percentages of 5.45%. In total, we identified 22 categories which can cover most of the resource-management APIs.

To further evaluate the security impact of these bugs, we follow the evaluation methods proposed by IPPO [25] and NDI [26]. Three types of security bugs are found by *ChatDetector*: memory leak, use-after-free and double-free. 95.7% of the detected bugs would cause memory leaks and lead to memory exhaustion or Denial-of-Services(DoS) when triggered repeatedly. 4 bugs would cause double-free and they could lead to memory corruption when triggered. One bug causes use-after-free and it may lead to DoS when triggered repeatedly. We then use a fixed bug as a case study of a common security impact of the detected bugs. In *HandBrake* with 15.2k stars in GitHub, the function `av_buffersrc_parameters_alloc()` allocates a new `AVBufferSrcParameters` instance and should be freed by the caller with `av_free()` according to the documentation. However, the parameter `par` is double freed both before and within the error-handling structure. Therefore, it will lead to a double-free bug. If the bug is triggered, *HandBrake* may encounter memory corruption. This bug is detected by *ChatDetector* and the bug is confirmed and fixed by the software holder and the latest version is linked below<sup>1</sup>.

### B. RQ2: Comparison with the state-of-the-art

To answer RQ2, we compare *ChatDetector* with the existing state-of-the-art vulnerability detection tools that utilize documentation analysis (Advance [2]) and code analysis (HERO [5] and Goshawk [4]) for RM-API misuse detection. Table II

<sup>1</sup><https://github.com/HandBrake/HandBrake/commit/ab467d13ef1f2a9a119c2cd241c7c831ca207192>

presents the detection results of these four approaches on three popular libraries: Libevent [10], Libexpat [11] and Libzip [13]. In total, *ChatDetector* detects 85 RM-API pairs and 11 security bugs on the three libraries. The other three tools (Advance, HERO and Goshawk) can only detect 9.41%, 1.17% and 7.05% of the RM-API pairs. All of them can merely detect 27.27%, 0% and 36.36% security bugs, respectively.

We further explore the reasons for the false negatives on RM-API misuse detection. Advance focuses on identifying sentences with strong sentiments, but most RM sentences do not have such tendencies, which leads to many misuses undiscovered. The difference between these two works is that *ChatDetector* can mine information beyond expert knowledge in identifying RM-API constraints. Although Advance can detect various API usage norms, the current RM APIs do not contain most of them and introduce many false negatives. HERO is designed to match functions accurately by utilizing the distinctive error-handling structure. The results show that HERO has limitations in capturing a large number of library functions, resulting from two primary factors. Firstly, the applications exhibit the absence of involving library functions, posing a challenge in extracting relevant library API usage from the code. Secondly, HERO’s emphasis on error-handling structures as a means to acquire API usages represents a unique approach not commonly employed in other software. Concerning Goshawk, designed to identify RM-API pairs through source code analysis, the lack of accurate API usage patterns can lead to significant false negatives in RM-API pairing and misuse detection. In this way, We found that *ChatDetector* outperforms other state-of-the-art tools in recognizing RM-API pairs and RM-API misuse detection.

```

1 import cpp
2 from Function f, AllocationFunction alloc,
  FunctionCall call
3 where call.getEnclosingFunction() = f and call.
  getTarget() = alloc
4 select f, f.getName()

```

Listing 2: QL code for detecting the memory leak of function `evconnlistener_new_bind`

Due to the static analysis tool CodeQL containing basic functionality for malloc API identification, we then compare *ChatDetector* against the basic functionality of CodeQL [27] named `AllocationFunction` and the QL code is shown in Listing 2. The results show that among all the six libraries, the basic function of CodeQL can only identify five allocation APIs which are `av_realloc` from `FFmpeg`, `pcap_list_datalinks`, `pcap_open_dead_with_tstamp_precision`, `pcap_list_tstamp_types` and `pcap_create` from `Libpcap`. The built-in malloc function identification functionality of CodeQL has difficulty in identifying most allocation APIs with multi-layered nesting. The five identified allocation APIs directly call `malloc` which can be easily identified by CodeQL. However, most of the allocation APIs have at least two-layered nesting indirectly calling

malloc and cannot be identified by CodeQL. For example, `evhttp_connection_base_bufferevent_new` from Libevent has allocation semantics and directly calls `mm_malloc`. Although `mm_malloc` directly calls `malloc` and represents clear allocation semantics, CodeQL cannot determine if `evhttp_connection_base_bufferevent_new` has allocation semantics. What’s more, `zip_source_buffer` from Libzip is composed of `malloc` through at least three-layered nesting and `pcap_findalldevs` from Libpcap is composed of `malloc` through at least four-layered nesting. Neither of them can be identified by CodeQL due to multi-layered nesting. The results show that the built-in malloc function identification of CodeQL conducts intra-procedural analysis which can only identify the allocation APIs directly calling `malloc` and leave a huge amount of RM-APIs with multi-layered nesting undiscovered (72.7% of the RM-APIs in GroundTruth have multi-layered nesting as illustrated in Section I). Therefore, the basic malloc function identification functionality of CodeQL is far less effective than *ChatDetector*, making it difficult to achieve the current task’s objectives.

We also evaluate *ChatDetector* on the standard API-misuse dataset APIMU4C [28], which contains 15 real-world RM-API misuses collected across three applications focusing on OpenSSL and basic C library. *ChatDetector* can detect 14 of them and all of them have the same bug type “memory leak”, the details are shown in Table III. The missing case highlighted in grey focuses on `X509_get_ext_d2i` which is deprecated with no direct detailed documentation and source code, therefore introduces false negatives of *ChatDetector*.

TABLE III: The RM-API misuses details from APIMU4C.

Library	Apps	API	Location
OpenSSL	httpd	X509_get_ext_d2i	httpd-4.37/modules/ssl/ssl_util_ssl.c:209
OpenSSL	httpd	BIO_new	httpd-4.37/modules/ssl/ssl_util_ssl.c:274
OpenSSL	httpd	BIO_new	httpd-4.37/modules/ssl/ssl_util_ocsp.c:361
OpenSSL	curl	BIO_new	curl-curl-7_63_0/lib/vtls/openssl.c:740
Basic C	curl	strdup	curl-curl-7_63_0/lib/vtls/openssl.c:1442
OpenSSL	curl	BIO_new	curl-curl-7_63_0/lib/vtls/openssl.c:2997
OpenSSL	curl	BIO_new	curl-curl-7_63_0/lib/vtls/openssl.c:3330
OpenSSL	curl	BIO_new	curl-curl-7_63_0/lib/vtls/openssl.c:3394
Basic C	curl	malloc	curl-curl-7_63_0/lib/escape.c:155
Basic C	curl	strdup	curl-curl-7_63_0/lib/smtp.c:518
Basic C	curl	aprintf	curl-curl-7_63_0/lib/smtp.c:531
Basic C	curl	strdup	curl-curl-7_63_0/lib/formdata.c:361
OpenSSL	openssl	OPENSSL_realloc	crypto/asn1/f_string.c:100
OpenSSL	openssl	EC_GROUP_dup	crypto/ec/ec_ameth.c:300
OpenSSL	openssl	EVP_MD_CTX_new	ssl/s3_enc.c:44

TABLE IV: The Results of Ablation Study

	Precision	Recall	F1 score
<i>Eval<sub>0</sub></i>	32.14%	29.67%	30.86%
<i>Eval<sub>1</sub></i>	77.97%	71.98%	74.9%
<i>ChatDetector</i>	98.21%	90.65%	94.28%

### C. RQ3: Ablation Study

In this section, we conduct two experiments among six libraries (FFmpeg [9], Libevent [10], Libexpat [11], Libpcap [12], Libzip [13] and OpenLdap [14]) to evaluate individual techniques. *ChatDetector* utilizes external knowledge, the Chain-of-Thought techniques and the cross-validation method for precisely RM-API pairing. Since the cross-validation should be implemented on the Chain-of-Thought results, we perform one evaluation for both of them. In the end, we evaluate the impact of different qualities of documentation for RM-API pairing. The details are shown as follows.

**Contribution of external knowledge.** To figure out the contribution of external knowledge for RM-API pairing, We design two types of prompts for RM sentence identification: **Type-1:** containing no documentation (directly using LLMs, denoting *Eval<sub>0</sub>* in Table IV); **Type-2:** containing API documentation (denoting *ChatDetector* in the same table). We further perform Chain-of-Thought and cross-validation with the identified RM sentences to RM-API pairing. Among the six libraries, the precision rate of *Eval<sub>0</sub>* on RM-API pairing is 32.14% and the F1 score is 30.86%, which are not nearly as effective as *ChatDetector*. The results fully explain the contribution of external knowledge to RM-API pairing.

**Contribution of Chain-of-Thought and cross-validation.** To evaluate the contribution of Chain-of-Thought and cross-validation, we also design two types of prompts: **Type-3:** prompting LLMs with external knowledge, denoting as *Eval<sub>1</sub>* in Table IV; **Type-4:** incorporating Chain-of-Thought and cross-validation with Type-3 prompting, denoting as *ChatDetector* in the same table. Performing Chain-of-Thought and cross-validation can obtain 25.94% more RM-API pairs with an increase in precision from 77.97% to 98.21% and an increase in recall from 71.98% to 90.65%, which demonstrates the contribution of Chain-of-Thought and cross-validation for RM-API pairing.

**The impact of documentation quality.** To evaluate the impact of different qualities of documentation for RM-API pairing, we design four types of prompts representing different qualities and incorporate them with *ChatDetector* to perform the evaluation: 1) directly prompting LLMs (referred to as No\_doc); 2) using the prompt that includes the first sentence of the description (referred to as First\_line); 3) using the prompt that includes the entire description (referred to as Full\_desc); 4) and using the prompt that includes the entire document (referred to as Full\_doc). The comparative results are shown in Figure 15. The results show that Full\_doc performs the best on the task among all the other prompts from three key dimensions: precision, recall and F1 score. Specifically, we




Fig. 15: The impact of different qualities of documentation on RM-API pairing

assume the API documentation with no description has the lowest quality and the API documentation with a one-sentence description also has significantly lower quality compared with the one including complete documentation. Figure 15 shows that the presence of documentation has a certain impact on RM-API pairing, as the F1 score rapidly increases when the description appears. However, as the document quality improves from a one-sentence description to a detailed description, the change in the F1 score is not significant, and the performance is clearly not as good as Full\_doc which is also applied in our work (*ChatDetector*). Therefore, we conclude that *ChatDetector* has limitations with undocumented API description, but does not show significant differences in other types of documentation quality.

## V. EXPLORATIONS AND CASE STUDY

Although the comparisons show that *ChatDetector* outperforms the state-of-the-art API misuse detectors, we are still curious about how *ChatDetector* achieves the task and to what extent it could be achieved. In this section, we further analyze the similarities and differences between *ChatDetector* and the benchmark method (human inspection) through three research questions. Additionally, we explore the potential future research directions for subsequent research.

**Benchmark study.** To mimic the whole human inspections of RM-API pairs identification through documentation analysis during the software development, including RM-API identification and RM-API pairing based on the same RM-object type, we chose three popular libraries for manual inspection evaluation: Libevent, Libzip and Libexpat. The complete manual inspection is illustrated in Figure 16. Take the function `zip_source_buffer` as an example, we first manually analyzed all the function calls within the function, through 6 iterations of function call analysis, we identified the core calling function `malloc` and confirmed the allocated object needs extra releasing operations, then we determined `zip_source_buffer` has allocation semantics, matching with the keyword “create” in the documentation. Adopting this method on all the functions from three libraries, we found six commonly-used keywords were used for describing RM semantics, including “allocate”, “create” and “construct” for

resource allocation semantics; “free”, “release” and “deallocate” for resource release semantics. In total, the GroundTruth contains 92 malloc APIs and their corresponding releasing APIs from three libraries, only 47 RM-API pairs can be identified through manual inspection.

**Exploration①: The relationship between buried API Constraints and the Documentation.** *ChatDetector* performs better than human inspection in mining constraints from API documentation. By analyzing the usage frequency of both existing and newly discovered resource-allocation functions in software, the statistics show that the newly discovered functions are not widely used in the software. Some functions only exist once, originating from the function’s source code rather than the calling code, indicating the rare usage identified in the development. The main difference between the benchmark method and *ChatDetector* is that humans identify allocation functions using documentation keywords, excluding the code-analysis method due to the complexity of the code. The results show that *ChatDetector* excels in mining RM-API pairs, identifying 0.8 times more RM-API pairs than the benchmark method. However, due to the relatively low usage frequency of identified RM-API pairs, the number of RM-API misuses detected by *ChatDetector* doesn’t significantly increase. *ChatDetector* discovered 38 new RM-API pairs, but only 34.21% of them are used, resulting in a total usage frequency of 36.63% compared to the benchmark, leading to detecting fewer misuses. Further, we infer the development cycle between the documentation writer, documentation and the users (referred to as software developers) through the analysis of the above phenomenon, as illustrated in Figure 17. When functions are infrequently used or issues are not immediately found, documentation depends on regular updates. Incomplete documentation can cause misuse. Frequently used or security-issued functions often get better documentation, causing imbalance. This imbalance may leave developers unaware of security guidelines for lesser-used functions, leading to security issues.

**Exploration②: How does *ChatDetector* identify RM sentences?** The statistical analysis of the identified RM-API pairs shows that *ChatDetector* introduces four resource-allocation APIs missing from identification due to the current method for extracting RM sentences. Most of the missed detections are attributed to the model itself. In the description, the RM sentence “*The zip\_open\_from\_source() function opens a zip archive ...*” accurately describes the functionality of `zip_open_from_source` and signifies that the current API has resource-allocation semantics. However, the model missed the identification and the misuse detection. Another reason for a missed detection is that the documentation itself lacks detailed descriptions, leaving the model without additional information to reference.

We conducted further analysis of RM sentences and found that the sentences containing keywords (as shown in the benchmark work, which is described in Section V) constitute 52.17% of them. The automation of the benchmark method for extracting RM sentences relies on matching corresponding keywords

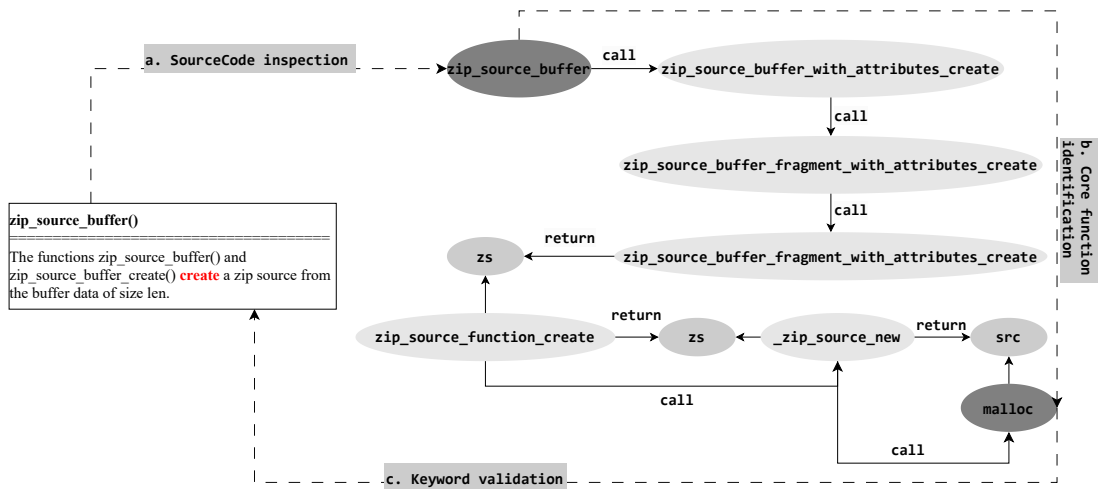


Fig. 16: Manual Inspection of zip\_source\_buffer

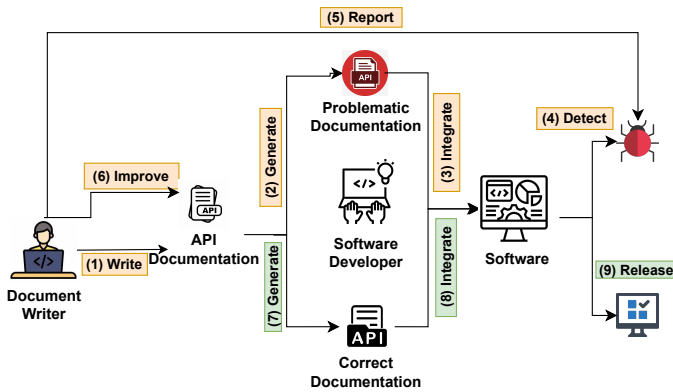


Fig. 17: The software development cycle

which may introduce both false positives and false negatives. However, thanks to *ChatDetector*'s comprehension abilities and the extensive training dataset, when identifying resource-allocation APIs, it goes beyond keyword-based recognition and can identify additional RM sentences. What's more, *ChatDetector* can detect RM sentences located in various parts of the documentation (e.g., description, first sentence, non-first sentence and return value) compared to traditional matching methods. The results reveal that leveraging the capabilities of LLMs could extend the boundaries of expert knowledge.

**Exploration⑧: How does *ChatDetector* identify RM APIs and paring them?** As a generative model, ChatGPT is trained on vast data and parameters, containing more expertise and a better ability to "understand". In contrast to the benchmark method, which is time-consuming and labour-intensive, *ChatDetector* acquires more knowledge and a well-designed solution for getting accurate answers from LLMs. For *RM-API Identification*, people should understand the API descriptions and further identify the RM APIs. But for *ChatDetector*, theoretically it should utilize the RM sentences as the evidence for RM-API identification. However, as a black-box model, it is hard to differentiate them from the training data. Therefore, *ChatDetector* can identify more RM sentences than the

**av\_get\_token()**

Unescape the given string until a non escaped terminating char, and return the token corresponding to the unescaped string.  
**Returns**  
 the malloced unescaped string, which must be av\_freed by the user, NULL in case of allocation failure

Fig. 18: The documentation defect found by *ChatDetector*

benchmark method. For *RM-API Pairing*, both the benchmark method and *ChatDetector* use a similar approach based on the same observation that resource-allocation functions and their corresponding releasing functions operate on objects of the same type. They achieve this by matching RM object descriptions with the function declarations of RM APIs to identify the RM object's type. The key difference lies in how they determine the RM-object type. For the benchmark method, people tend to calculate the semantic similarity between RM object descriptions and the names of parameter/return value types to recognize the RM object type. In contrast, *ChatDetector* simplifies this step into one single prompt. For most incomplete or loosely-organized descriptions, *ChatDetector* demonstrates stronger text-processing capabilities and can accurately identify RM-object types with relatively less information. What's more, Figure 18 shows a defect found by *ChatDetector* in the documentation. People are supposed to pair *av\_get\_token* with *av\_freed* according to the description, which means the caller of *av\_get\_token* is supposed to call *av\_freed* to prevent potential memory leak. However, the current presentation of the documentation is not only hard to follow, but also mentions the incorrect function, which does not exist. Such types of defects can result in incorrect constraints being extracted from existing work [2]. However, *ChatDetector* is not only unaffected by incorrect information, but also capable of detecting such defects within the documentation.

*ChatDetector* uses the RM-object type provided in the prompt to determine the corresponding releasing function,



```

ZIP_EXTERN int
zip_set_archive_comment(zip_t *za, const char
*comment, zip_uint16_t len) {
    zip_string_t *cstr;

    if (ZIP_IS_RDONLY(za)) {
        zip_error_set(&za->error, ZIP_ER_RDONLY, 0);
        return -1;
    }

    if (len > 0 && comment == NULL) {
        zip_error_set(&za->error, ZIP_ER_INVAL, 0);
        return -1;
    }

    if (len > 0) {
        if ((cstr = _zip_string_new((const zip_uint8_t
*)comment,
len, ZIP_FL_ENC_GUESS, &za->error)) == NULL)
            return -1;
        ...
    }
}

```

Fig. 19: The source code of `zip_set_archive_comment`

while the benchmark method splits the procedure into RM object type identification and corresponding release API identification. In terms of results, *ChatDetector* can discover more RM-API pairs and achieve higher accuracy. However, since ChatGPT is a black-box model and its training data is not accessible, it may produce judgments indicating resource allocation semantics even when there is no clear evidence of RM semantics in the text. For example, ChatGPT’s analysis of the documentation for the function `zip_set_archive_comment` led it to conclude that the function has resource allocation semantics. Manual inspection of the source code in Figure 19 reveals that the function source code calls `_zip_string_new` to allocate the resource of the string `cstr` and then promptly frees it by calling `_zip_string_free` within the function. Since this string is not passed as the return value or a variable to the function `zip_set_archive_comment`, the function indeed has resource-allocation semantics but does not require further resource releasing. Therefore, *ChatDetector*’s false positive for this function might be due to its internal dataset containing the source code of this function. When *ChatDetector* prompts LLMs for an analysis of the function `zip_set_archive_comment`, the response from LLMs may also include an analysis of the source code, leading to *ChatDetector*’s false positive.

## VI. DISCUSSIONS

**The use of other techniques.** As mentioned in Section III, we apply the widely-used GPT-3.5-turbo model [22] to implement *ChatDetector*. The total cost of six libraries using GPT-3.5-turbo is cheaper than 8.9 dollars and the cost of static program analysis is within a reasonable price. We also compared GPT-4.0-turbo and ChatGPT, not only do the hallucinations still exist, but the cost of GPT-4.0-turbo is higher including the monthly user cost and the 20 times increased expense on each query. Traditional type analysis [29], [30] needs full source-code analysis of the function to acquire the type, while the LLM only takes one refined prompt for RM-object type

identification which reaches nearly 93.55% precision and the cost is cheaper than 0.00058 dollars per function on average. Therefore, utilizing LLMs exceeds the traditional type analysis in type identification.

**Limitations.** *ChatDetector* introduces the novel concept of utilizing LLMs to understand and validate its capability in analyzing official API documentation and extracting RM-API constraints beyond expertise for RM-API misuse detection. However, due to the black-box nature of LLMs, it is challenging to understand how the model generates certain false positives and addresses them equally. We also conducted a small comparison between CodeLLama-7b [31], vicuna-7b [32], GPT-3.5-turbo [22] and GPT-4.0 [33] shows that hallucination still exists across different LLMs where directly using them will generate incorrect answers for RM-API identification. Considering their effectiveness and the corresponding costs, we chose GPT-3.5 to implement *ChatDetector*. *ChatDetector* attempts to enhance the stability of LLMs’s output inspired by ReAct framework, but this issue cannot be entirely eliminated, which remains further discussion.

**Future Work.** With the increasing capabilities of LLMs, the future of text understanding and analysis research has shifted from breaking through the bottlenecks of traditional NLP tools to exploring how to effectively leverage LLMs to assist in other tasks. As tools make breakthrough advancements, tasks that were previously challenging due to tool limitations will present new opportunities. From one side, since the current explainable AI methods have not yet reached the field of LLMs [34], [35], [36], the interpretability of LLMs could be a new direction worth exploring, which helps in exploring the potential of LLMs. From the other side, building upon the foundation of this work, the scope of extracted constraints can be gradually expanded, and further analysis of complex documentation can be conducted to enhance the work of API misuse detection.

## VII. RELATED WORK

**Resource-Management API misuse detection.** Current static vulnerability detection methods can be divided into two types of approaches: code-based analysis and documentation-based analysis. Previous research assumed API constraints were implemented in code and detected misuse by recovering constraints. Some works assume releasing resources post-allocation through analysis of error-handling structures [5], [37]; other works suggest mining frequent function usages to detect misuses [6], [38]. However, code-based methods assume successful API rule implementation, leading to false negatives and positives due to pattern extraction. Previous studies extracted API constraints from documentation to aid in detecting misuse. Some work identifies Resource Management (RM) APIs from Javadoc using resource-related keywords [1]; other works involving templates matching within organized documentation [39] and extracting Integrated Assumptions (IAs) from loosely organized C documentation using sentiment-emphasized constraints [2]. Our work innovatively utilizes



LLMs for documentation understanding, adopting the inspirations from the ReAct framework, identifying the resource-allocation function, along with its resource-releasing function and constructing the RM-API constraints.

**Large Language Models in Security Research.** In the era of LLMs' widespread adoption, research on their use as alternatives for addressing security issues has emerged. Some researchers propose to utilize LLMs for fuzzing: Deng et al [40], [41] proposed the direct utilization of LLMs to generate input programs, enabling effective fuzz testing of deep learning libraries. Another use is automated bug fixing using LLMs. Both Pearce et al.[42] and Xia et al.[43], [44] propose LLM-based bug fixing methods, confirming safe repairs without extra fine-tuning. However, the current research poses a blank space for utilizing LLMs for API misuse detection. Our work identifies concealed API usage patterns by comprehending the content of API documentation. It verifies the authenticity and accuracy of the answers through further processing and analysis of the reasoning process applied to LLMs.

### VIII. CONCLUSION

In this paper, we introduce an automated RM-API misuse detector, *ChatDetector*, through task decomposition inspired by the ReAct framework to understand API official documentation with LLMs, extracting RM-API constraints with a two-dimensional prompting for cross-validation and inconsistency-checking between the LLMs' output and reasoning process, with which enables the real-world RM-API misuse detection. In total, *ChatDetector* discovers 115 security issues, that could potentially leading severe security issues, and are ethically reported to the developers. Compared to the state-of-the-art API misuse detectors, *ChatDetector* can additionally detect at least 80.85% more RM-API pairs with a 98.21% precision, demonstrating the ability of LLMs on information retrieval beyond expertise. This work enlightens new research directions for API misuse detection aided by LLMs, extending beyond overcoming the bottlenecks of traditional NLP tools to explore how to effectively leverage LLMs for security research.

### ACKNOWLEDGEMENTS

We want to thank our shepherd and reviewers for their insightful and constructive comments which highly improve our paper. The IIE authors are supported in part by NSFC (U24A20236, 92270204), CAS Project for Young Scientists in Basic Research (Grant No. YSBR-118), Youth Innovation Promotion Association CAS and the Open Research Fund of Joint Laboratory on Cyberspace Security, China Southern Power Grid (Grant NO. GDKJXM20230642).

### REFERENCES

[1] H. Zhong, L. Zhang, T. Xie, and H. Mei, "Inferring specifications for resources from natural language api documentation," *Automated Software Engineering*, vol. 18, no. 3, pp. 227–261, 2011.

[2] T. Lv, R. Li, Y. Yang, K. Chen, X. Liao, X. Wang, P. Hu, and L. Xing, "Rtfm! automatic assumption discovery and verification derivation from library document for api misuse detection," in *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, 2020, pp. 1837–1852.

[3] R. Pandita, X. Xiao, H. Zhong, T. Xie, S. Oney, and A. Paradkar, "Inferring method specifications from natural language api descriptions," in *2012 34th international conference on software engineering (ICSE)*. IEEE, 2012, pp. 815–825.

[4] Y. Lyu, Y. Fang, Y. Zhang, Q. Sun, S. Ma, E. Bertino, K. Lu, and J. Li, "Goshawk: Hunting memory corruptions via structure-aware and object-centric memory operation synopsis," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 2096–2113.

[5] Q. Wu, A. Pakki, N. Emamdoost, S. McCamant, and K. Lu, "Understanding and detecting disordered error handling with precise function pairing," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 2041–2058.

[6] I. Yun, C. Min, X. Si, Y. Jang, T. Kim, and M. Naik, "Apsan: Sanitizing api usages through semantic cross-checking," in *UseNix Security Symposium*, 2016, pp. 363–378.

[7] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao, "React: Synergizing reasoning and acting in language models," 2023.

[8] "Semantic role labeling," [https://web.stanford.edu/~jurafsky/slp3/slides/22\\_SRL.pdf](https://web.stanford.edu/~jurafsky/slp3/slides/22_SRL.pdf), 2023.

[9] "Ffmeg documentation," <http://ffmpeg.org/>, 2023.

[10] "Libevent documentation," <https://libevent.org/doc/>, 2023.

[11] "Libexpat documentation," <https://libexpat.github.io/doc/api/latest/>, 2023.

[12] "Libpcap documentation," <https://www.tcpdump.org/manpages/pcap.3pcap.html>, 2023.

[13] "Libzip documentation," <https://libzip.org/documentation/libzip.html>, 2023.

[14] "Openldap documentation," <https://www.openldap.org/software/man.cgi?query=ldap>, 2023.

[15] S. Feng and C. Chen, "Prompting is all your need: Automated android bug replay with large language models," *arXiv preprint arXiv:2306.01987*, 2023.

[16] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell et al., "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.

[17] 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 preprint arXiv:2303.12712*, 2023.

[18] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou et al., "Chain-of-thought prompting elicits reasoning in large language models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 824–24 837, 2022.

[19] "Codeql," <https://codeql.github.com/>, 2023.

[20] J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus, "Emergent abilities of large language models," 2022.

[21] G. Agrawal, T. Kumarage, Z. Alghami, and H. Liu, "Can knowledge graphs reduce hallucinations in llms? : A survey," 2023.

[22] "Openai documentation," <https://platform.openai.com>, 2023.

[23] J. Maynez, S. Narayan, B. Bohnet, and R. McDonald, "On faithfulness and factuality in abstractive summarization," 2020.

[24] C. Parsing, "Speech and language processing," *Power Point Slides*, 2009.

[25] D. Liu, Q. Wu, S. Ji, K. Lu, Z. Liu, J. Chen, and Q. He, "Detecting missed security operations through differential checking of object-based similar paths," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 1627–1644.

[26] Q. Zhou, Q. Wu, D. Liu, S. Ji, and K. Lu, "Non-distinguishable inconsistencies as a deterministic oracle for detecting security bugs," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 3253–3267.

[27] "Native codeql," [https://codeql.github.com/codeql-standard-libraries/cpp/semml/code/cpp/models/interfaces/Allocation.qll/type.Allocation\\$AllocationFunction.html](https://codeql.github.com/codeql-standard-libraries/cpp/semml/code/cpp/models/interfaces/Allocation.qll/type.Allocation$AllocationFunction.html), 2024.

[28] Z. Gu, J. Wu, J. Liu, M. Zhou, and M. Gu, "An empirical study on api-misuse bugs in open-source c programs," in *2019 IEEE 43rd annual computer software and applications conference (COMPSAC)*, vol. 1. IEEE, 2019, pp. 11–20.

[29] Y. Chen, Z. Lin, and X. Xing, "A systematic study of elastic objects in kernel exploitation," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020, pp. 1165–1184.

- [30] K. Lu, A. Pakki, and Q. Wu, “Detecting Missing-Check Bugs via Semantic- and Context-Aware Criticalness and Constraints Inferences,” in *Proceedings of the 28th USENIX Security Symposium (Security)*, Santa Clara, CA, 2019.
- [31] “Codellama-7b,” <https://huggingface.co/samavr/codellma-7b>, 2024.
- [32] “Vicuna-7b,” <https://huggingface.co/lmsys/vicuna-7b-v1.5>, 2024.
- [33] “Gpt-4.0,” <https://openai.com/index/gpt-4/>, 2024.
- [34] A. Holzinger, A. Saranti, C. Molnar, P. Biecek, and W. Samek, “Explainable ai methods-a brief overview,” in *International workshop on extending explainable AI beyond deep models and classifiers*. Springer, 2022, pp. 13–38.
- [35] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, “Explainable ai: A review of machine learning interpretability methods,” *Entropy*, vol. 23, no. 1, p. 18, 2020.
- [36] F. Xu, H. Uszkoreit, Y. Du, W. Fan, D. Zhao, and J. Zhu, “Explainable ai: A brief survey on history, research areas, approaches and challenges,” in *Natural language processing and Chinese computing: 8th cCF international conference, NLPCC 2019, dunhuang, China, October 9–14, 2019, proceedings, part II 8*. Springer, 2019, pp. 563–574.
- [37] S. Saha, J.-P. Lozi, G. Thomas, J. L. Lawall, and G. Muller, “Hector: Detecting resource-release omission faults in error-handling code for systems software,” in *2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2013, pp. 1–12.
- [38] P. Bian, B. Liang, J. Huang, W. Shi, X. Wang, and J. Zhang, “Sinkfinder: harvesting hundreds of unknown interesting function pairs with just one seed,” in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2020, pp. 1101–1113.
- [39] R. Pandita, K. Taneja, L. Williams, and T. Tung, “Icon: Inferring temporal constraints from natural language api descriptions,” in *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2016, pp. 378–388.
- [40] 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*, 2023, pp. 423–435.
- [41] Y. Deng, C. S. Xia, C. Yang, S. D. Zhang, S. Yang, and L. Zhang, “Large language models are edge-case fuzzers: Testing deep learning libraries via fuzzgpt,” *arXiv preprint arXiv:2304.02014*, 2023.
- [42] H. Pearce, B. Tan, B. Ahmad, R. Karri, and B. Dolan-Gavitt, “Examining zero-shot vulnerability repair with large language models,” in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2022, pp. 1–18.
- [43] C. S. Xia, Y. Wei, and L. Zhang, “Automated program repair in the era of large pre-trained language models,” in *Proceedings of the 45th International Conference on Software Engineering (ICSE 2023)*. Association for Computing Machinery, 2023.
- [44] C. S. Xia and L. Zhang, “Keep the conversation going: Fixing 162 out of 337 bugs for \$0.42 each using chatgpt,” *arXiv preprint arXiv:2304.00385*, 2023.

## IX. APPENDIX

### A. Details of evaluation

Several APIs are used for illustrating the effectiveness of *ChatDetector*, for example the documentation of `zip_source_begin_write_cloning` shown in Figure 20 and `zip_open_from_source` shown in Figure 21.

The functions `zip_source_begin_write()` and `zip_source_begin_write_cloning()` prepare source for writing. **Usually this involves creating temporary files or allocating buffers.** `zip_source_begin_write_cloning()` preserves the first offset bytes of the original file. This is done efficiently, and writes to source won't overwrite the original data until `zip_commit_write()` is called.

Fig. 20: The documentation of `zip_source_begin_write_cloning`

The `zip_open()` function opens the zip archive specified by path and returns a pointer to a struct `zip`, used to manipulate the archive. The flags are specified by or'ing the following values, or 0 for none of them.

`ZIP_CHECKCONS`  
Perform additional stricter consistency checks on the archive, and error if they fail.

`ZIP_CREATE`  
Create the archive if it does not exist.

`ZIP_EXCL`  
Error if archive already exists.

`ZIP_TRUNCATE`  
If archive exists, ignore its current contents. In other words, handle it the same way as an empty archive.

`ZIP_RDONLY`  
Open archive in read-only mode. If an error occurs and `errorp` is non-NULL, it will be set to the corresponding error code.

**The `zip_open_from_source()` function opens a zip archive encapsulated by the `zip_source` `zs` using the provided flags. In case of error, the `zip_error` `ze` is filled in.**

Fig. 21: The documentation of `zip_open_from_source`

We choose six popular libraries for *ChatDetector* implementation and evaluation (Section IV), the details are shown in Table V.

TABLE V: Library Information

Library	Function Description	#Line of Doc sentences	#APIs	#Apps
Ffmpeg	Multimedia	3238	734	645
Openldap	Lightweight Directory Access Protocol	5854	126	239
Libevent	Asynchronous event notification	2175	401	85
Libpcap	Provides a portable framework for low-level network monitoring	454	71	181
Libzip	Reading, creating, and modifying zip archives	2416	120	27
Libexpat	Parsing XML	422	70	275
<b>Total</b>	/	<b>14,559</b>	<b>1,522</b>	/

TABLE VI: Discovered RM-API misuses

Lib	Software	Allocation API	Releasing API	Bug Type	#Bugs
libevent	seafile	event_base_new	event_base_free	memleak	1
	evpp	event_base_new	event_base_free	memleak	2
		event_config_new	event_config_free	memleak	1
		event_new	event_free	memleak	2
		evhttp_new	evhttp_free	memleak	1
transmission	event_new	event_free	memleak	1	
libzip	radare2	zip_source_buffer_create	zip_source_free	memleak	1
	OpenRCT2	zip_source_buffer	zip_source_free	memleak	1
ffmpeg	gpac	av_dict_set	av_dict_free	memleak	2
		av_dict_copy	av_dict_free	memleak	1
	HandBrake	av_dict_set	av_dict_free	memleak	21
		avfilter_graph_create_filter	avfilter_free	memleak	1
		av_buffersrc_parameters_alloc	av_free	double free	1
	FFmpeg	av_frame_new_side_data	av_frame_remove_side_data	memory leak	29
		av_frame_alloc	av_frame_free	double free	1
		av_new_program	av_free	memory leak	8
	owntone-server	av_dict_set	av_dict_free	memleak	3
	vlc	av_malloc	av_freep	memleak	1
libpcap	PF_RING	pcap_findalldevs	pcap_freealldevs	memleak	2
		pcap_compile	pcap_freecode	memleak	4
	arp-scan	pcap_compile	pcap_freecode	memleak	1
	freeradius-server	pcap_compile	pcap_freecode	memleak	1
	nmap	pcap_compile	pcap_freecode	memleak	1
	ntopng	pcap_compile	pcap_freecode	memleak	2
	tcpdump	pcap_compile	pcap_freecode	memleak	1
	wireshark	pcap_compile	pcap_freecode	memleak	2
	openvas-scanner	pcap_compile	pcap_freecode	memleak	1
pcap_findalldevs		pcap_freealldevs	memleak	2	
ldap	freebsd-src	ldap_search_s	ldap_msgfree	memleak	3
	freeradius-server	ldap_result	ldap_msgfree	double free	2
		ldap_result	ldap_msgfree	memleak	1
	gpdb	ldap_search_s	ldap_msgfree	memleak	1
	openldap	ldap_search_ext_s	ldap_msgfree	memleak	6
		ldap_url_parse	ldap_free_urldesc	memleak	2
		ldap_initialize	ldap_unbind_ext	use after free	1
		ldap_initialize	ldap_unbind_ext	memleak	2
yugabyte-db	ldap_search_s	ldap_msgfree	memleak	1	
postgres	ldap_search_s	ldap_msgfree	memleak	1	
Total	-	-	-	-	115