

# REEF: A Framework for Collecting Real-World Vulnerabilities and Fixes

Chaozheng Wang<sup>§¶</sup>, Zongjie Li<sup>‡¶</sup>, Yun Peng<sup>†</sup>, Shuzheng Gao<sup>†</sup>, Sirong Chen<sup>§</sup>  
Shuai Wang<sup>‡</sup>, Cuiyun Gao<sup>§\*</sup>, Michael R. Lyu<sup>†</sup>

<sup>§</sup> School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China

<sup>‡</sup> Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong, China

<sup>†</sup> Computer Science and Engineering Department, The Chinese University of Hong Kong, Hong Kong, China

{wangchaozheng, 190110916}@stu.hit.edu.cn, {zligo, shuaiw}@cse.ust.hk

{ypeng, szgao23, lyu}@cse.cuhk.edu.hk, gaocuiyun@hit.edu.cn

**Abstract**—Software plays a crucial role in our daily lives, and therefore the quality and security of software systems have become increasingly important. However, vulnerabilities in software still pose a significant threat, as they can have serious consequences. Recent advances in automated program repair have sought to automatically detect and fix bugs using data-driven techniques. Sophisticated deep learning methods have been applied to this area and have achieved promising results. However, existing benchmarks for training and evaluating these techniques remain limited, as they tend to focus on a single programming language and have relatively small datasets. Moreover, many benchmarks tend to be outdated and lack diversity, focusing on a specific codebase. Worse still, the quality of bug explanations in existing datasets is low, as they typically use imprecise and uninformative commit messages as explanations.

To address these issues, we propose an automated collecting framework REEF to collect REal-world vulnERabilities and Fixes from open-source repositories. We focus on vulnerabilities since they are exploitable and have serious consequences. We develop a multi-language crawler to collect vulnerabilities and their fixes, and design metrics to filter for high-quality vulnerability-fix pairs. Furthermore, we propose a neural language model-based approach to generate high-quality vulnerability explanations, which is key to producing informative fix messages. Through extensive experiments, we demonstrate that our approach can collect high-quality vulnerability-fix pairs and generate strong explanations. The dataset we collect contains 4,466 CVEs with 30,987 patches (including 236 CWE) across 7 programming languages with detailed related information, which is superior to existing benchmarks in scale, coverage, and quality. Evaluations by human experts further confirm that our framework produces high-quality vulnerability explanations.

**Index Terms**—Vulnerability, Data collection, Bug fix

## I. INTRODUCTION

Software powers an increasing number of critical systems in the modern world, from people's daily life applications [1] to critical industrial scenarios [2] or even military systems [3]. However, software often contains bugs that can cause substantial losses [4]. In this paper, we focus on vulnerabilities since they are a more severe type of bug related to software security that may have serious consequences. The Common Vulnerabilities and Exposures (CVE) database tracks publicly disclosed cybersecurity vulnerabilities and the number

of reported CVEs has been growing rapidly over the past decade [5], highlighting the scale and significance of this problem; this growth is largely attributed to the increasing number of open-source software applications, as well as the increasing sophistication of cyber-attacks and vulnerability discovery methods. To combat the potential risk in real-world software, various approaches have been proposed to detect and fix bugs.

Traditionally, static analysis tools have been used to analyze source code and detect potential bugs. A large number of mature tools such as Coverity [6] have been maintained and improved for several decades, applying to real-world projects to discover potential flaws. More recently, query-based code search tools like CodeQL [7] and Sengrep [8] have enabled developers to detect bugs by writing semantic queries to explore code bases. The core idea behind the tools is to provide basic analysis capabilities to developers and let them write queries to detect bugs. The effectiveness of these tools depends on the quality of the queries [9], which keeps the detection process both flexible and scalable.

Besides the traditional static analysis and query-based code search tools, data-driven approaches [10]–[19] are also promising for detecting and fixing bugs. Different from analyzing the source code with predefined rules along with complex analysis techniques (e.g., dataflow analysis and control flow analysis), data-driven approaches leverage large-scale data from open-source communities such as GitHub to learn the patterns of bugs and fixes. Techniques such as zero-day detector [20] generate candidate patches and validate them to find viable fixes. Other data-driven methods learn developer-written patches from open-source repositories and apply them to repair new bugs.

The advantages of the data-driven approaches are two-fold. First, they are easier to implement than traditional static analysis tools. The traditional tools require complex analysis techniques to analyze the source code, which is time-consuming and hard to scale. For example, when facing a wide variety of software vulnerabilities, professional developers must carefully design appropriate detection rules or techniques [6], [7]. In contrast, the data-driven approaches only require large-scale data from open-source repositories, which is easy to obtain

<sup>¶</sup> These authors contribute equally.

\* Cuiyun Gao is the corresponding author.

and scale. Second, they are more flexible than the traditional tools. The traditional tools are usually language-agnostic and designed for specific types of bugs, which are hard to extend to other programming languages or other types of bugs [21]–[23]. The data-driven approaches are more flexible, as they can be easily extended to other types of bugs by learning the patterns of bugs and fixes from the large-scale data.

Data-driven approaches for automated bug detection and repair rely on the availability of large-scale, high-quality datasets. The efficacy of these techniques is directly dependent on the precision and comprehensiveness of the data used to train machine learning models. Specifically, accurately identifying bug locations and patches can bolster the precision of the trained models, improving the accuracy of bug detection and fixing. Furthermore, incorporating metadata on bug types and commit messages provides models with more granular information about bugs, which can enhance their performance on detection and repair tasks [24]. This additional context may also aid developers during debugging and code review by giving them more targeted insights.

However, current datasets for data-driven software analysis have several key limitations that hamper their effectiveness. First, the granularity of most datasets is at the function level, lacking precise records of bug locations and fixes. In reality, bugs often span multiple levels of abstraction and a single CVE may impact various, disparate parts of a codebase [25]. Thus, there is a shortage of systematically curated, real-world vulnerability data. Second, metadata about bug types is often inaccurate or imprecise. Currently, bug types are typically inferred from commit messages, which can be inaccurate and even sometimes wrong [26]. Such an error-prone process would possibly yield incorrect or misleading labels. Third, most current datasets [27] are outdated, failing to capture the latest state of constantly evolving software systems. Bugs that were previously patched may be reintroduced, tending to render existing records obsolete.

**Technical Challenges and Solutions.** We aim to develop a framework for automatically collecting and curating high-quality code snippets containing vulnerabilities, fixes, locations, their types, and messages from open-source repositories. It is thus required to prepare a large dataset incorporating this information to gain insights into real-world bugs and fixes, facilitating further research and applications. To achieve the above goals, our approach comprises three steps: ❶ **Newest CVE capture:** current datasets focus on simple and outdated vulnerability patterns without clear location information as well as the patch. To develop a comprehensive dataset that is close to real-world situations, we primarily focus on the recently revealed CVEs that share clear fix log and location information. This results in a dataset with about eighteen thousand CVEs that are recently revealed with documented fixes. ❷ **Large Language Model (LLM)-based explanation with human agreement:** Current datasets adopt the commit information as the comment for the vulnerability content, which is proved to be unreliable [28], [29]. Considering the powerful capability of LLM in code understanding [30],

[31], we leverage the large language model to automatically comprehend the bug patterns and fixes from the CVEs and use them as the additional message. To ensure the quality of the mined results, we carefully design the prompt system with a pilot study. Moreover, we conduct a human agreement study to evaluate the mined results. ❸ **Analysis of the collected dataset:** we conduct further analysis of the dataset to understand the characteristics of real-world bugs and fixes, providing details on each data point to benefit future studies or applications.

Our contributions are summarized as follows:

- We propose REEF, a framework to mine up-to-date, real-world vulnerabilities automatically. Incorporated with the corresponding fix patches from CVEs, we have collected a large-scale dataset with 30,987 bug location, type, and fix information. Our dataset consists of a wide variety of vulnerabilities across various languages, platforms, and granularity.
- We employ large language models to generate explanatory messages for the CVEs, supplementing unreliable commit information. We carefully design prompts and conduct a human evaluation to ensure message quality.
- We conduct an extensive analysis of the collected dataset and provide detailed insights into real-world vulnerabilities and fix characteristics to guide future research and applications.
- We have publicly released the code for our REEF tool on GitHub at <https://github.com/ASE-REEF/REEF-script>, along with the vulnerability data we have collected, which is also available on GitHub at <https://github.com/ASE-REEF/REEF-data>.

## II. RELATED WORK

In this section, we introduce the related work from three threads including automated program repair, static analysis tools, and large language models for code, respectively.

### A. Automated Program Repair

Automatic program repair (APR) has garnered significant attention in recent years as a crucial approach to enhancing software reliability. Various techniques have been developed within the APR domain, including template-based [32], [33], search-based [34], [35], constraint-based [36], [37], and learning-based approaches [10]–[19]. Among these categories, learning-based methods have achieved the greatest success and become the most popular in recent years. SequenceR [10] combines LSTM encoder-decoder architecture with copy mechanism for program repair. DLFix [13] uses a tree-based RNN to capture the structure of the source code and learn code transformations. CoCoNuT [14] uses ensemble learning on the combination of different networks to automatically fix bugs in multiple programming languages, separating the context and buggy lines in NMT-based APR. CURE [15] integrates pre-trained programming language models and significantly improves repair quality. Recoder [11] uses a syntax-guided edit decoder to guide the generation of syntactically correct repair patches. RewardRepair [16] employs execution-based

backpropagation to enhance the compilation rate of patches generated by NMT-based APR approaches. DEAR [17] generate multi-hunk, multi-statement fix patches with a divide-and-conquer strategy. AlphaRepair [12] utilizes a large pre-trained code model and generates patches in a fill-in-the-blank way. Zhong et al. [38] build a standard benchmark dataset and an extensive framework tool to mitigate threats for comparison in program repair. Xia et al. [39] evaluate the effectiveness of LLMs on program repair. KNOD [18] incorporates domain knowledge to guide patch generation in a direct and comprehensive way. TypeFix [19] is a prompt-based approach with fix templates incorporated for repairing Python type errors.

### B. Static Analysis Tools

Static analysis tools analyze source code without executing the program to detect potential bugs and vulnerabilities. They codify definitions of unsafe coding patterns and scan codebases to identify matches. Over the past decade, static analysis has become a popular approach for detecting vulnerabilities in software [40], and techniques including data flow analysis [41], [42], tpestate analysis [43], type inference [44] and specific pointer analysis [45] have been developed to improve the precision and recall of bug detection.

Recently, query-based static analysis tools like CodeQL [7] have gained increasing attention from both academia and industry. These tools codify vulnerability patterns as SQL-like queries, facilitating knowledge sharing and reuse across entities and software systems. Software is treated as data [46], with programs parsed into hierarchical representations, often stored in databases. Unlike traditional static analysis tools [6], [47], query-based tools primarily focus on parsing software into rich, query-friendly representations and rely on crowdsourced communities to continually develop queries addressing newly discovered vulnerabilities. Established query-based tools cultivate active communities and offer bounty programs [48] to encourage query contribution and improvement. In turn, these communities help enrich and refine queries to target vulnerabilities proliferating in real-world software.

There is a huge effort to establish comprehensive benchmarks to evaluate the quality of analysis tools, which further helps find real-world vulnerabilities. However, existing benchmarks for evaluating static analysis tools typically use synthetic datasets. For instance, the Juliet [21] benchmark for C/C++ and the Defects4J [22] in Java language. These follow prescribed patterns and quickly become outdated, unable to represent the complexity of real-world CVEs. Although built around the Common Weakness Enumeration (CWE) to provide reasonable, well-defined examples, they cannot capture the nuances of most vulnerabilities. Some works collect datasets from real CVEs and make complex processes to filter the suitable ones. For instance, Ruohonen [49] targets Python and collects samples from popular repositories and Linares et al. [50] analyze Android apps. These narrow scopes limit the types of vulnerabilities and coding patterns addressed, impeding holistic analysis. However, these works usually focus

on particular programming languages or codebases, lacking diversity.

While static analysis shows promise for detecting vulnerabilities at scale, evaluating tools remains challenging due to the lack of comprehensive benchmarks reflecting the diversity of real-world bugs. Real CVEs offer a rich source for dataset generation but are difficult and time-consuming to gather and curate. Progress in static analysis thus depends on developing datasets that mirror the heterogeneity of vulnerabilities in real code. Automated or semi-automated methods for collecting and labeling examples from a wide range of open-source repositories directly show potential for advancing research and practice in this crucial area of software security, which is the focus of our work.

### C. Large Language Models for Code

Recently, significant advancements in SE research have been brought by Large Language Models (LLMs), which brought impressive improvements in a wide range of code-related tasks. One notable model is InCoder [51], which employs a causal masking training objective to excel in code infilling and synthesis. Another popular model is Codex [52], a sizable pre-trained code model introduced by OpenAI, which supports the Copilot service on various code-related tasks [53]. The models recently released by OpenAI, such as ChatGPT [54] and GPT-4 [55], are also pre-trained with source code data and show remarkable programming capabilities. AlphaCode [56] has been specifically trained for generating code for programming competitions like Codeforces. CodeCMR [57] and IRGEN [58] are pre-trained models designed for low-level code on various code-related tasks. CodeGen [59] is a large pre-trained model for multi-turn program synthesis with more than 16B parameters, while CodeGeeX [60] is a recently proposed open-source multilingual code generation model with 13 billion parameters. BigCode Project has developed and open sourced StarCoder [61] which contains 15.5B parameter. A recent work WizardCoder [62] is fine-tuned with Evol-Instruct and achieves state-of-the-art performance surpassing all existing open-source Code LLMs.

In-context learning (ICL) [63], [64] is a recent paradigm that enables LLMs to learn from just a few examples without fine-tuning. It concatenates a few input-output examples with the query question to form an input for the language model and get the prediction. Recently, there has been increasing interest in applying in-context learning to code-related tasks [39], [65]–[67]. CEDAR [66] retrieves similar examples and constructs the demonstrations for assert generation and program repair. Synchromesh [68] retrieves few-shot examples by Target Similarity Tuning and samples programs using Constrained Semantic Decoding. A recent work [69] empirically studies the impact of three demonstration construction factors on in-context learning in code intelligence tasks. Geng et al. [70] enhance in-context learning for multi-intent code comment generation by selecting similar examples and re-ranking the output candidates. Ahmed et al. [71] propose to incorporate

static analysis results into the in-context prompt to code summarization.

### III. WORKFLOW

Fig. 1 presents our pipeline for collecting real-world vulnerabilities and constructing the dataset. Our pipeline consists of three steps: ❶ CVE capturing and collection. To address the limitations of current benchmarks, we gather recent real-world vulnerabilities from CVEs that have been newly disclosed, including related bug reports and source code. We use an automated crawler to collect the latest CVEs and then filter out irrelevant ones based on metrics like CVSS scores. ❷ LLM message supplementation. To account for uneven commit message quality and potential bias, we leverage the ability of large language models [72] to generate vulnerability explanations for each commit. Following standard bug description guidelines [73], we design a prompt for the vulnerability explanations using advanced LLM, enabling the model to formulate consistent, unified messages through contextual learning. ❸ Dataset analysis. We introduce metrics to assess our dataset's quality and compare it with current benchmarks to determine effectiveness in supporting existing tools. We analyze the generated messages and compare them with committed information to evaluate our approach. Finally, we conduct human studies to assess the quality of the generated messages.

To evaluate the effectiveness of our approach to collect code vulnerability and the quality of our proposed dataset, we investigate the following three research questions (RQs):

- RQ1:** What is the advantage of our dataset compared to existing benchmarks?
- RQ2:** To what extent the prompt design affects the generated message?
- RQ3:** How are the generated messages in alignment with experts?

Specifically, we analyze our collected dataset and compare it with current benchmarks to explore whether it is effective and diverse in RQ1 (in Section IV-A). As we use the code understanding ability of LLM, in RQ2 we further study how the prompt would affect the performance of generated bug explanation compared to the commit information in Section IV-B. Finally in RQ3, by comparing the patches we collected as well as the generated message, we study to what extent humans are in agreement with the generated code explanation as discussed in Section IV-C.

Notably, we focus on a wide variety of vulnerabilities across various languages and platforms, not only including the commonly-seen pattern bugs occur at the function level or statement level, but also the complex vulnerabilities across multiple files and functions, which we believe is more challenging for the current tools to detect and can better benefit the community.

#### A. Data Collection

Creating an exhaustive dataset including real-world code snippets with vulnerabilities, fixes, locations, and types is challenging, let alone including vulnerability explanations.

To develop a dataset reflecting real-world scenarios, we first gather real-world vulnerabilities from multiple sources, including the NVD database and CVE list maintained by Mend [74], which is a comprehensive open-source vulnerabilities database from hundreds of both popular and under-the-radar community resources. Users can also specify additional sources as needed. If a vulnerability has a clear report, proof-of-concept, and publicly available source code before and after fixing, we collect the related bug reports and source code, and store them in our raw dataset. Notably, it would be possible that a single vulnerability may be linked to multiple commits and files; we gather all related commits and files accordingly.

As shown in Fig. 1, we design a filter to remove less severe vulnerabilities from the raw dataset. We first eliminate those with a low CVSS score which indicates relatively low impact and damage potential. We then assign a “fix score” based on the number of related commits and a weighted score to each commit based on the number of files modified. Vulnerabilities with low “fix scores” are excluded from the final dataset.

To expand the dataset's potential applications and adapt it for various downstream tasks, we incorporate disclosure date information for each item. Specifically, we keep the complete CVE name indicating when each vulnerability was disclosed and assigned a unique number. Users preferring more recent data can easily filter out vulnerabilities disclosed during a given time period.

#### B. LLM Message Supplementation

To construct a comprehensive, informative dataset, vulnerability descriptions are crucial since they provide rationale and fix details, helping downstream tools better understand vulnerabilities. This is especially useful for data-driven methods, as natural language is easier to comprehend than code [75]. However, commit messages are not always available and may be uninformative, biased [76], or misleading [26] in explaining vulnerabilities. Worse yet, the existing commit messages could be missing, impairing downstream approaches. To address this, we leverage large language models to generate vulnerability explanations for each commit. Following standard bug description guidelines [73], we empirically design a prompt for the vulnerability explanations using advanced LLM, enabling the model to formulate consistent, unified messages through in-context learning.

After collecting the responses from the large language model, we further conduct a human inspection to evaluate whether the generated message is in alignment with experts and whether it is informative enough to explain the vulnerability. Only the generated message that is in agreement with the experts would be included in the final dataset.

We provide a complete list of field names for our collected dataset in Table I. The fields can be categorized into four groups: (1) metadata, (2) vulnerability information, (3) LLM-enhanced information, and (4) project information. The metadata describes the programming language and index for each data item. The vulnerability information contains CVE details documenting the real-world impact of the bugs. The

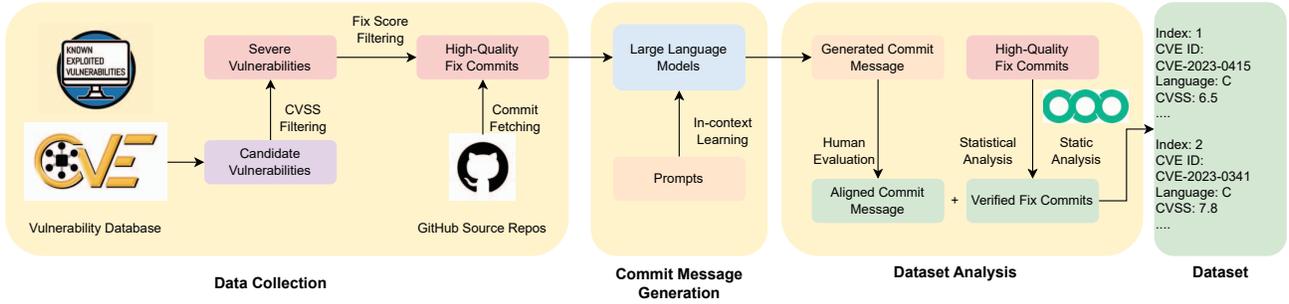


Fig. 1. The pipeline of our REEF for gathering vulnerabilities, enriching data, and analyzing the dataset.

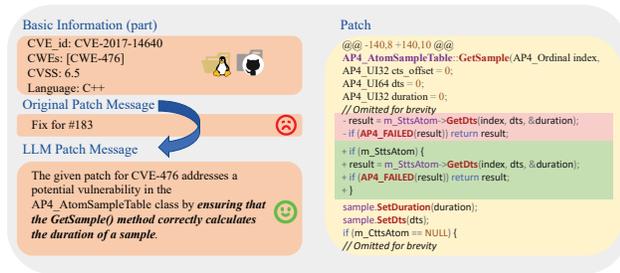


Fig. 2. Example of data instance with enhanced patch messages.

LLM-enhanced messages are generated descriptions of the vulnerabilities.

The project information stores data on the code repositories mined for bugs, including website details. In summary, our dataset incorporates comprehensive metadata, security details, generated explanatory messages, and provenance information to enable in-depth analysis of real-world vulnerabilities. The multi-faceted data provides insights into vulnerability properties, remediation, language model performance, and codebase characteristics.

### C. Dataset Analysis

To rigorously evaluate our dataset’s quality, we conduct an exhaustive analysis. Conceptually, we compare our dataset to current benchmarks across several axes: supported programming languages, fixing information integrity, granularity, data source, corpus size, and CWE coverage, respectively. A comparative analysis along these dimensions provides a holistic sense of relative strengths. A systematic assessment of vulnerability type diversity examines coverage of the Common Weakness Enumeration (CWE). We posit that a dataset exhibiting a wider range of CWE types affords more comprehensive vulnerability modeling, with greater potential to generalize across systems. By comparing the number of CWE types, we gain quantitative insight into our dataset’s diversity. We further evaluate the challenge to current tools using Semgrep [8], a ubiquitous query-based static analysis tool, to detect vulnerabilities in our dataset. A higher proportion of vulnerabilities evading detection suggests greater resilience against existing methods, indicating the dataset poses a more

TABLE I  
DESCRIPTION OF FIELDS IN THE JSON FILE OF OUR DATASET.

Field Name in JSON	Description
<b>Metadata</b>	
index	The index of the instance in the dataset
language	The programming language of the instance
<b>Vulnerability Information</b>	
cve_id	The identifier of CVE that the instance belongs to
cvss	The CVSS score reflecting the severity of the CVE
cwes	The CWEs, i.e., the common weaknesses in the CVE
<b>LLM-enhanced Information</b>	
llm_message	The commit message generated by LLMs
<b>Project Information</b>	
origin_message	The original commit message
url	The URL of the commit on GitHub REST API
html_url	The URL that locates the webpage of the commit
raw_url	The URL of the source file changed in the commit
raw_code	The source code of the file changed in the commit

formidable evaluation benchmark. Missed vulnerabilities point to remaining gaps in vulnerability modeling and detection that the dataset could help address through continued research. Finally, we conduct an in-depth statistical analysis of our dataset itself to glean qualitative details that inform future work. Summary statistics on parameters like vulnerabilities’ average severity and exploitability, bug fix length, and explanation shine a light on real-world characteristics underrepresented in synthetic data. A granular exploration of attributes exposes new problem dimensions beyond the capabilities of simplified synthetic benchmarks.

These complementary analyses, systematically connected through a logical pipeline, provide empirical evidence and qualitative characterization to demonstrate our dataset’s diversity, challenge, and fidelity in emulating real-world scenarios. The rigor and depth of our evaluative approach underline the dataset’s potential to serve as a foundation for future research advancing the state of the art in software vulnerability detection and automatic program repair.

## IV. EVALUATION

*A. RQ1: What is the advantage of our dataset compared to existing benchmarks?*

**Conceptual comparison.** We first conduct a conceptual comparison between our selected dataset and current datasets, in which we compare the differences during the dataset collection process, including the multi-language support, fix information, location, related message, granularity, source, and size.

As shown in Table II, our dataset has the following advantages compared to current benchmarks: (1) Detailed fix information. Compared to other benchmarks that only include fix patterns, our dataset contains detailed fix information, including the bug location and fix information. (2) Multi-language support. Different from the existing datasets that focus on specific languages, ours incorporate vulnerabilities in multiple languages, including C/C++, Java, C#, and Python. (3) Multi-level granularity. Our dataset contains vulnerabilities at multiple levels, including function-level, statement-level, and expression-level. (4) Real-world CVE. The dataset we collected contains real-world CVEs, which are more representative than synthetic vulnerabilities. (5) Large-scale. A vast volume of 30,987 vulnerability patches enables robust statistics and enhances machine learning via increased instances, far more than other benchmarks' limited samples.

In summary, at the conceptual level, our dataset is more comprehensive and representative than existing ones.

**Dataset coverage.** As discussed in Sec. III-C, we further analyze the coverage of Common Weakness Enumeration (CWE) types in our dataset compared to other benchmarks. Our hypothesis is that a dataset exhibiting a wider range of CWE types will enable more comprehensive vulnerability modeling with greater potential for generalization across systems.

As shown in Table II, our dataset covers more CWE types than all other benchmarks. Due to the limitation of space, we only show the total number of all CWE types across languages, but our dataset covers more CWE types even when focusing on one specific language. For example, our dataset covers 134 CWE types in C/C++, while Juliet-C++ benchmark only covers 118 CWE types. Notably, it would be hard to estimate a clear number of potential CWEs in a specific language, since some of the CWEs are not language-specific. However, we can still observe that our dataset covers more CWE types than other benchmarks, which indicates that our dataset provides more comprehensive vulnerability coverage compared to existing benchmarks.

Beyond CWE coverage analysis, we also evaluate detected CWE coverage using a static analysis tool. The intuition is that a lower proportion of successfully detected vulnerabilities suggests a more challenging dataset, as existing flaws are harder to discover with standard tools. Such difficulty highlights the potential utility of advanced models. Specifically, we use Semgrep [8], a popular query-based static analysis tool, to detect potential vulnerabilities. We use its default ruleset, which contains 1,088 rules for Java, 655 rules for Python, and 133 rules for JS, to detect vulnerabilities in our

dataset and other benchmarks. Defects4J only provides code patches, which makes analyzing it difficult. Consequently, we do not compare our benchmark with Defects4J. Juliet-C++ has a notably high detection rate at 35.7%, presumably because it is a synthetic dataset. LinuxFlaw and FUNDED, the benchmarks containing real-world CVEs, have substantially lower detection rates of 2.2% and 4.0%, respectively. In contrast, our dataset has an even lower detection rate of 1.2%, indicating that our dataset poses a more challenging problem. This is likely due to the fact that our dataset comprises more recent and intricate real-world vulnerabilities mined from up-to-date CVEs. Consequently, defects in our dataset may be more difficult to be detected using current approaches.

In summary, our collected dataset covers more CWE types and is more challenging for standard static analysis tools to assess, suggesting it is more comprehensive and representative than existing benchmarks. The broader range of hard-to-detect vulnerabilities in our dataset could support more robust vulnerability modeling and lead to repair systems with stronger generalization ability.

**Dataset statistics.** We first analyze the statistics of our collected dataset. As shown in Table III, our dataset contains 4466 vulnerabilities across seven languages. For each vulnerability, we report the average number of changed files, patch numbers, and changed lines of code, with results shown in the table. These results demonstrate that Java and C# vulnerabilities are substantially more complex than those in the other languages, with average values for all metrics nearly double those of the other languages. This aligns with our expectation that these are usually served as Object-Oriented (OO) languages, and appear more compact in the coding specification. Moreover, we observe that C language still consists of a large proportion of vulnerabilities, which indicates that (1) C language is still widely used in practice. (2) The unfamiliarity with the C language leads to ongoing vulnerability discovery. As an old language compared to Python, the language abstraction of C is relatively low, and developers need to manage memory manually (e.g. buffer overflows and memory leaks). This feature makes C language more error-prone, which leads to a large number of vulnerabilities.

Beyond analysis at the source code level, we also examine our dataset in terms of vulnerability types. As shown in Fig. 3, the top 15 CWE types constitute over 55% of our dataset, indicating it is comprehensive and representative. We also observe that Java and Python have similar CWE type distributions, while C/C++ has a distinct distribution. The top five CWE types are CWE-79 (Cross-Site Scripting), CWE-125 (Out-of-bounds Read), CWE-787 (Out-of-bounds Write), CWE-20 (Improper Input Validation), and CWE-119 (Improper Restriction of Operations within the Bounds of a Memory Buffer). It is unsurprising that the dataset over-represents these CWE types, which are archetypical examples of memory-related vulnerabilities. Overall, these statistics suggest that our dataset effectively samples from the space of vulnerabilities and contains a diversity of complexity levels and types, especially for C/C++ and low-level languages.

TABLE II  
COMPARISON OUR COLLECTED DATASET AND CURRENT BENCHMARKS ✓, ◊, ✗ DENOTE SUPPORT, PARTIALLY SUPPORT, AND NOT SUPPORT, RESPECTIVELY.

	# Sup language	Fix information	location	Related Msg	Granularity	Source	Size	# CWE Types
Juliet-C++ [21]	1	✗	✓	✗	statement-level	Synthetic	64,099	118
Defects4J [22]	1	✓	✗	✗	function-level	Synthetic	835	-
LinuxFlaw [23]	1	✓	✗	✓	multi-level	CVE	368	18
FUNDED [77]	4	✓	✗	✗	function-level	Synthetic & CVE	6561	54
Ours	7	✓	✓	✓	multi-level	CVE	30,987	236

TABLE III  
STATISTICS OF OUR COLLECTED DATASET.

Languages	# Case	# Func	# Avg diff file	# Avg patch	# Avg COL
C++	411	2244	2.88	5.46	86.81
C	1575	6957	2.14	4.42	62.97
Java	541	6207	5.74	11.47	297.13
Python	863	5797	3.26	6.72	113.2
JS	636	5066	4.26	7.97	130.32
Go	355	3187	4.54	8.98	195.43
C#	85	1529	8.98	17.99	201.29
Total	4466	30987	4.54	9.0	155.30

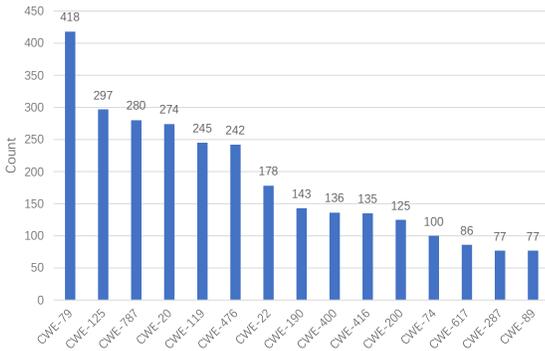


Fig. 3. Top-15 CWE types of our dataset.

**Finding 1:** REEF, the framework we propose to collect the code vulnerabilities, considers more aspects than existing benchmarks, which in turn makes the dataset more practical to the real world. Through comprehensive analysis, we demonstrate that our dataset covers more CWE types and is more challenging than other benchmarks, indicating its comprehensiveness and representativity than other benchmarks.

*B. RQ2: To what extent does the prompt design affect the generated message?*

We answer RQ2 from two aspects. First, we conduct a pilot study on C/C++ projects to analyze the generated message from the perspective of their basic understanding of vulnerabilities and fix records. Second, we generate the message by using three commonly-used prompt patterns and compare them with the original message.

**Pilot Study.** Inspired by [78], we list the investigated LLMs in Table IV. We use GPT-Neo [79], Llama-7B, and Llama-13B [80] fine-tuned on Alpaca [81] and released by LM-Flow [82]. Vicuna is Llama-based and fine-tuned on user-

TABLE IV  
LLMs USED IN THE EXPERIMENT, WITH PLAUSIBILITY COUNTS FOR “COMPLETELY TRACEABLE”, “SOMEWHAT TRACEABLE” AND “NON-TRACEABLE” LABELS. THE SYMBOL “-” DENOTES THAT THE CORRESPONDING STATISTIC IS UNKNOWN.

Model	Vendor	Year	# Para.	Traceability
GPT-Neo [79]	EleutherAI	2021	2.7B	5/3/7
Llama-7B [80]	Meta	2023	7B	5/6/4
Llama-13B [80]	Meta	2023	13B	7/4/4
ChatGLM [83]	THU	2023	6B	10/4/1
Vicuna [84]	BAIR	2023	13B	13/0/2
ChatGPT [54]	OpenAI	2022	-	13/1/1
Tongyi [85]	Alibaba	2023	-	13/1/1
GPT-4 [55]	OpenAI	2023	-	15/0/0

shared conversations. ChatGLM [83] and Vicuna [84] models use official code. We also include commercial LLMs ChatGPT [54] and GPT-4 [55]. To make a fair comparison, we use the same prompt for all models. Moreover, we randomly select 15 CVE cases from C/C++ vulnerabilities. After generating the message, we manually label the message as “completely traceable”, “somewhat traceable” and “non-traceable”. Where if the message incorporates the information in the CVE patch or the vulnerability rationale, we label it as “completely traceable”; if the message incorporates the information in the CVE patch or the vulnerability rationale partially, we label it as “somewhat traceable”; otherwise, we label it as “non-traceable”.

From the results, we observe that ChatGPT and Tongyi gain competitive performance in generating messages that are traceable to the CVE patch or the vulnerability rationale besides GPT-4. Considering the limitations of APIs’ accessibility and academic resources, we use Tongyi as the baseline model in the following experiments.

**Prompt pattern comparison.** Notably, we only generate one summarization message for each CVE, regardless of the number of commits and files. It could be the situation that one CVE contains multiple commits and files, which means the generated message is a summarization of all the commits and files. However, since the input token length is limited, we set the maximum number of generated explanation tokens to 256. Once the input surpasses the limit, we truncate the input and generate the message based on the truncated input. Users can increase this number to generate more detailed explanations if needed.

Besides the token length, the prompt pattern is also a key point. It has been demonstrated that the output performance for code-related tasks is largely influenced by the type of prompt

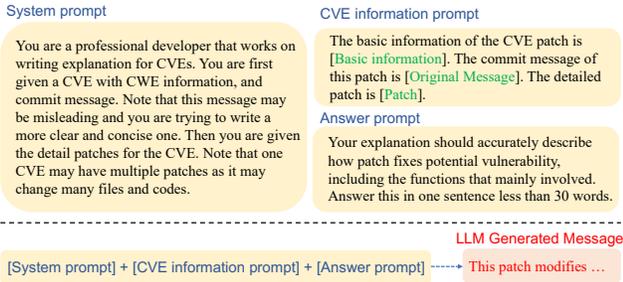


Fig. 4. Example of the prompt we used to query large language models.

pattern [86]. Specifically, we select 20 cases to generate messages based on the three prompt patterns, named “Zero-shot”, “One-shot” and “Few-shot”.

The prompts used to generate patch messages in the zero-shot setting are presented in Figure 4. For one-shot and few-shot prompts, we give the large language model one or a few human-authored examples to guide the model in generating high-quality patch messages. Here we do not consider the Chain-of-thought [87] pattern since no suitable rationale for code is available. For each prompt pattern, we generate the message based on the CVE patch and the vulnerability rationale. We then invite five experts to evaluate the generated message based on the following criteria:

- **Comprehensiveness:** whether the message is comprehensive enough to explain the vulnerability and fix.
- **Consistency:** whether the message is consistent with the CVE patch and the vulnerability rationale.
- **Traceability:** whether the message is traceable to the CVE patch and the vulnerability rationale.

Each expert is asked to give a score range of 0 to 1 to demonstrate the generated quality regarding the corresponding criteria. After collecting the evaluation results, we calculate the average score for each prompt pattern. The results are shown in Table V. From the results, we observe that the “Few-shot” pattern achieves the highest score in all three criteria, while the “One-shot” pattern shares competitive results. As a trade-off between the performance and computation cost, we adopt “One-shot” as the default prompt pattern in this work.

TABLE V  
PROMPT PATTERN COMPARISON.

Prompt Pattern	Comprehensiveness	Consistency	Traceability
Zero-shot	0.65	0.7	0.6
One-shot	1.0	0.95	0.95
Few-shot	1.0	1.0	0.95

**Finding 2:** Large language models differ in their ability to generate vulnerability explanations traceable to source code and rationale. Tongyi and GPT models excelled in a pilot study, while a “one-shot learning” prompt achieved a satisfying performance for generating comprehensive, consistent, and traceable explanations in a prompt pattern comparison with an affordable query budget.

*C. RQ3: How are generated messages in alignment with experts?*

To answer RQ3, we first present statistics comparing our generated messages to the original commit messages. We then conduct a human evaluation to assess the quality of our generated messages relative to the corresponding commit messages.

**Statistics comparison** We report statistics for our generated messages and the original commit messages in Table VI. On average, our generated messages contain 397.08 characters, 1.92 times more than the original commit messages. Moreover, the median length of our generated messages is 356, also higher than for the original commit messages. We attribute this to the fact that original commit messages are written by developers with varying perspectives and goals, yielding greater diversity than our generated messages.

Notably, some original commit messages are of low quality for two reasons: (1) They are auto-filled by the GitHub platform. (2) They are too short (less than 20 characters.) We count the number of these low-quality commit messages and report them in the “Lcmsg” column of Table III. Of the original commit messages in our dataset, nearly 5% of messages are viewed as of low quality. They are unsuitable for providing informative vulnerability explanations and were thus excluded from our analyses, further motivating our work to generate comprehensive, high-quality vulnerability explanations.

The statistics demonstrate that our generated messages are substantially more detailed and consistent than the original commit messages. The higher word count suggests our messages provide more structured and in-depth vulnerability explanations overall compared to original commit messages, which are often quite brief and arbitrary. The filtering of low-quality messages is also prudent, as their inclusion could skew the statistics and make the dataset unsuitable for modeling. These results thus indicate we achieve our aim of generating high-quality, comprehensive vulnerability explanations.

**Human Study.** We conduct a human evaluation to assess the quality of the generated messages. We recruit five experts, including two industrial developers and three academic researchers with expertise in software vulnerability detection, as participants. We randomly select 40 samples and create an online questionnaire for them. For each sample, we provide two messages without specifying their origin. On a scale of 1 to 5 (1 being completely unsatisfactory, 5 being fully satisfactory), participants score each message. To ensure participants understand the task, we include five sanity check (SC) test items, considering only participants who answer all SC items correctly. Each participant evaluates 35 real samples and five

TABLE VI  
STATISTICS OF OUR GENERATED MESSAGE AND THE ORIGINAL COMMIT MESSAGE.

Languages	# Case	# Lcmsg	# Acmsg (Med.)	# Agmsg (Med.)
C++	411	21	234.93 (156)	415.02 (364)
C	1,575	122	380.0 (148)	389.78 (351)
Java	541	38	152.63 (68)	399.51 (356)
Python	863	36	204.11 (125)	408.19 (363)
JS	636	60	123.74 (57.0)	382.84 (346.0)
Go	355	20	237.68 (86)	401.15 (376)
C#	85	3	109.85 (52)	383.13 (340)
Total	4,466	300	206.13 (98.85)	397.08 (356.57)

SC items; we assign each sample to five participants. All participants pass the SC, taking an average of 35 minutes.

The human evaluation finds an average score of 3.05 for original messages and 3.70 for generated messages, a 21.31% relative gain. Analysis of all responses shows that for 7.14% of cases, the generated message seems worse, while for the remaining 92.86%, the generated message is equal to or better than the original. The Fleiss' Kappa [88] of 0.92 indicates "almost perfect agreement" between participants. These results suggest the generated messages are well-aligned with expert assessments and higher in quality than the original commit messages. The small minority of cases where the generated message seems inferior could be anomalous or reflects subtle aspects not captured in our message-generation approach. But the overwhelming expert preference for the generated messages, further evidenced by strong inter-rater agreement, indicates their quality is superior overall.

**Finding 3:** Our generated vulnerability explanations are of superior quality to original commit messages according to both expert human evaluations and quantitative message statistics. Though a small fraction of generated messages are inferior, experts overwhelmingly prefer our generated explanations, indicating they are well-aligned with human assessments of explanation quality.

## V. DISCUSSION

**Limitations and Threats to Validity.** We now discuss the validity and limitations of this work. In this research, we collect the dataset from the real-world CVEs where the open-source projects are hosted on GitHub. However, not all open-source projects are hosted on GitHub, and we may miss some projects that are hosted on other platforms. Moreover, since we only collect the CVEs from 2016, the previous CVEs are not included in our dataset. This may cause bias in the dataset, as some of the critical vulnerabilities may be discovered and fixed in previous years. In addition, we only collect the CVEs that are fixed by the developers. However, some CVEs are not fixed by the developers, or the developers have already fixed the vulnerabilities but did not report or confirm the CVEs. In summary, our collection framework tried to collect the CVEs as much as possible, but it is still possible that some CVEs are missing in our dataset, where future work can improve.

**Message bias.** In this work, we leverage the code understanding ability of large language models to generate additional messages for the collected vulnerabilities. Although we have specified strict rules as well as a dedicated designed prompt system, and our human evaluation shows that the generated messages are of high quality, some of the generated messages may still be biased or unsatisfactory. For example, when the input code snippet is either too long or too short, the generated message may not be as good as the other cases. However, their superior ratings compared to original commit messages, as assessed by experts, indicate they still achieve the aim of producing informative vulnerability explanations, even if imperfect. Moreover, the quality of generated messages can be further improved with the help of more advanced language models.

## VI. CONCLUSION

In this paper, we propose a novel and practical approach to collect the real-world CVEs with detailed information automatically, and we leverage the code understanding ability of large language models to generate additional messages for the collected vulnerabilities. Our framework successfully collects 4,466 CVEs from 2016 to 2023 and incorporates 30,987 messages for the collected patches. The collected dataset has been evaluated by the developers, and the results show that the generated messages are of high quality and can help the developers to understand the vulnerabilities. This work serves as a roadmap for researchers to construct better data-driven bug detection and auto-fix techniques.

## ACKNOWLEDGMENT

This research is supported by the National Natural Science Foundation of China under project (No. 62002084), Natural Science Foundation of Guangdong Province (Project No. 2023A1515011959), Shenzhen Basic Research (General Project No. JCYJ20220531095214031 and No. JCYJ20200109113403826). Additionally, this research is supported by the Key-Area Research and Development Program of Guangdong Province (No. 2020B010165002), and is also supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14206921 of the General Research Fund).

## REFERENCES

- [1] A. Jadhav, M. Kaur, and F. Akter, "Evolution of software development effort and cost estimation techniques: five decades study using automated text mining approach," *Mathematical Problems in Engineering*, vol. 2022, pp. 1–17, 2022.
- [2] G. Li, Y. Hou, and A. Wu, "Fourth industrial revolution: technological drivers, impacts and coping methods," *Chinese Geographical Science*, vol. 27, pp. 626–637, 2017.
- [3] W. Wang, H. Liu, W. Lin, Y. Chen, and J.-A. Yang, "Investigation on works and military applications of artificial intelligence," *IEEE Access*, vol. 8, pp. 131 614–131 625, 2020.
- [4] R. Telang and S. Wattal, "An empirical analysis of the impact of software vulnerability announcements on firm stock price," *IEEE Transactions on Software Engineering*, vol. 33, no. 8, pp. 544–557, 2007.
- [5] "Mitre," <https://cve.mitre.org/>.
- [6] Synopsys, "Coverity," <https://scan.coverity.com/>, 2022.

- [7] “CodeQL Introduction,” <https://help.semml.com/QL/ql-training/cpp/intro-ql-cpp.html>.
- [8] r2c, “Semgrep,” <https://semgrep.dev>, 2021.
- [9] R. Yang, J. Cai, and X. Han, “Poster: Taintgrep: A static analysis tool for detecting vulnerabilities of android apps supporting user-defined rules,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 3507–3509.
- [10] Z. Chen, S. Kommrusch, M. Tufano, L. Pouchet, D. Poshyanyk, and M. Monperrus, “Sequencer: Sequence-to-sequence learning for end-to-end program repair,” *IEEE Trans. Software Eng.*, vol. 47, no. 9, pp. 1943–1959, 2021.
- [11] Q. Zhu, Z. Sun, Y. Xiao, W. Zhang, K. Yuan, Y. Xiong, and L. Zhang, “A syntax-guided edit decoder for neural program repair,” in *ESEC/FSE ’21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Athens, Greece, August 23–28, 2021*. ACM, pp. 341–353.
- [12] C. S. Xia and L. Zhang, “Less training, more repairing please: revisiting automated program repair via zero-shot learning,” in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2022, Singapore, Singapore, November 14–18, 2022*. ACM, 2022, pp. 959–971.
- [13] Y. Li, S. Wang, and T. N. Nguyen, “Dlfix: context-based code transformation learning for automated program repair,” in *ICSE ’20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*. ACM, 2020, pp. 602–614.
- [14] T. Lutellier, H. V. Pham, L. Pang, Y. Li, M. Wei, and L. Tan, “Coconut: combining context-aware neural translation models using ensemble for program repair,” in *ISSTA ’20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, USA, July 18–22, 2020*. ACM, 2020, pp. 101–114.
- [15] N. Jiang, T. Lutellier, and L. Tan, “CURE: code-aware neural machine translation for automatic program repair,” in *43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22–30 May 2021*. IEEE, 2021, pp. 1161–1173.
- [16] H. Ye, M. Martinez, and M. Monperrus, “Neural program repair with execution-based backpropagation,” in *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25–27, 2022*. ACM, 2022, pp. 1506–1518.
- [17] Y. Li, S. Wang, and T. N. Nguyen, “DEAR: A novel deep learning-based approach for automated program repair,” in *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25–27, 2022*. ACM, 2022, pp. 511–523.
- [18] N. Jiang, T. Lutellier, Y. Lou, L. Tan, D. Goldwasser, and X. Zhang, “KNOD: domain knowledge distilled tree decoder for automated program repair,” *CoRR*, vol. abs/2302.01857, 2023.
- [19] Y. Peng, S. Gao, C. Gao, Y. Huo, and M. R. Lyu, “Domain knowledge matters: Improving prompts with fix templates for repairing python type errors,” *CoRR*, vol. abs/2306.01394, 2023.
- [20] W. Cui, M. Peinado, H. J. Wang, and M. E. Locasto, “Shieldgen: Automatic data patch generation for unknown vulnerabilities with informed probing,” in *2007 IEEE Symposium on Security and Privacy (SP’07)*. IEEE, 2007, pp. 252–266.
- [21] nist, “Juliet Test Suite for C/C++,” <https://samate.nist.gov/SRD/testsuite.php>, 2021.
- [22] R. Just, D. Jalali, and M. D. Ernst, “Defects4j: A database of existing faults to enable controlled testing studies for java programs,” in *Proceedings of the 2014 international symposium on software testing and analysis*, 2014, pp. 437–440.
- [23] D. Mu, A. Cuevas, L. Yang, H. Hu, X. Xing, B. Mao, and G. Wang, “Understanding the reproducibility of crowd-reported security vulnerabilities,” in *27th {USENIX} security symposium ({USENIX} security 18)*, 2018, pp. 919–936.
- [24] J. Bai, L. Zhou, A. Blanco, S. Liu, F. Wei, M. Zhou, and Z. Li, “Jointly learning to repair code and generate commit message,” in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 9784–9795.
- [25] A. Nappa, R. Johnson, L. Bilge, J. Caballero, and T. Dumitras, “The attack of the clones: A study of the impact of shared code on vulnerability patching,” in *2015 IEEE symposium on security and privacy*. IEEE, 2015, pp. 692–708.
- [26] Z. Liu, X. Xia, A. E. Hassan, D. Lo, Z. Xing, and X. Wang, “Neural-machine-translation-based commit message generation: how far are we?” in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 373–384.
- [27] Y. Zhou, S. Liu, J. Siow, X. Du, and Y. Liu, “Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks,” *Advances in neural information processing systems*, vol. 32, 2019.
- [28] A. Bachmann, C. Bird, F. Rahman, P. Devanbu, and A. Bernstein, “The missing links: bugs and bug-fix commits,” in *Proceedings of the eighteenth ACM SIGSOFT international symposium on Foundations of software engineering*, 2010, pp. 97–106.
- [29] Y. Tian, Y. Zhang, K.-J. Stol, L. Jiang, and H. Liu, “What makes a good commit message?” in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 2389–2401.
- [30] Y. Peng, S. Gao, C. Gao, Y. Huo, and M. R. Lyu, “Domain knowledge matters: Improving prompts with fix templates for repairing python type errors,” *arXiv preprint arXiv:2306.01394*, 2023.
- [31] S. Gao, X.-C. Wen, C. Gao, W. Wang, and M. R. Lyu, “Constructing effective in-context demonstration for code intelligence tasks: An empirical study,” *arXiv preprint arXiv:2304.07575*, 2023.
- [32] K. Liu, A. Koyuncu, D. Kim, and T. F. Bissyandé, “Tbar: revisiting template-based automated program repair,” in *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2019, Beijing, China, July 15–19, 2019*. ACM, 2019, pp. 31–42.
- [33] A. Koyuncu, K. Liu, T. F. Bissyandé, D. Kim, J. Klein, M. Monperrus, and Y. L. Traon, “Fixminer: Mining relevant fix patterns for automated program repair,” *Empir. Softw. Eng.*, vol. 25, no. 3, pp. 1980–2024, 2020.
- [34] J. Jiang, Y. Xiong, H. Zhang, Q. Gao, and X. Chen, “Shaping program repair space with existing patches and similar code,” in *Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2018, Amsterdam, The Netherlands, July 16–21, 2018*. ACM, 2018, pp. 298–309.
- [35] C. L. Goues, T. Nguyen, S. Forrest, and W. Weimer, “Genprog: A generic method for automatic software repair,” *IEEE Trans. Software Eng.*, vol. 38, no. 1, pp. 54–72, 2012.
- [36] Y. Pei, Y. Wei, C. A. Furia, M. Nordio, and B. Meyer, “Code-based automated program fixing,” in *26th IEEE/ACM International Conference on Automated Software Engineering (ASE 2011), Lawrence, KS, USA, November 6–10, 2011*. IEEE Computer Society, 2011, pp. 392–395.
- [37] Y. Pei, C. A. Furia, M. Nordio, Y. Wei, B. Meyer, and A. Zeller, “Automated fixing of programs with contracts,” *IEEE Trans. Software Eng.*, vol. 40, no. 5, pp. 427–449, 2014.
- [38] W. Zhong, H. Ge, H. Ai, C. Li, K. Liu, J. Ge, and B. Luo, “Standup4npr: Standardizing setup for empirically comparing neural program repair systems,” in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 2022, pp. 1–13.
- [39] C. S. Xia, Y. Wei, and L. Zhang, “Practical program repair in the era of large pre-trained language models,” *CoRR*, vol. abs/2210.14179, 2022.
- [40] B. Chess and G. McGraw, “Static analysis for security,” *IEEE security & privacy*, vol. 2, no. 6, pp. 76–79, 2004.
- [41] E. Bodden, “Inter-procedural data-flow analysis with ifds/ide and soot,” in *Proceedings of the ACM SIGPLAN International Workshop on State of the Art in Java Program analysis*, 2012, pp. 3–8.
- [42] Z. Liu, D. Xiao, Z. Li, S. Wang, and W. Meng, “Exploring missed optimizations in webassembly optimizers,” in *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2023, pp. 436–448.
- [43] S. J. Fink, E. Yahav, N. Dor, G. Ramalingam, and E. Geay, “Effective typesafe verification in the presence of aliasing,” *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 17, no. 2, pp. 1–34, 2008.
- [44] Y. Peng, C. Gao, Z. Li, B. Gao, D. Lo, Q. Zhang, and M. Lyu, “Static inference meets deep learning: a hybrid type inference approach for python,” in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 2019–2030.
- [45] Y. Smaragdakis, G. Balatsouras et al., “Pointer analysis,” *Foundations and Trends® in Programming Languages*, vol. 2, no. 1, pp. 1–69, 2015.
- [46] F. Yamaguchi, N. Golde, D. Arp, and K. Rieck, “Modeling and discovering vulnerabilities with code property graphs,” in *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 590–604.
- [47] META, “Facebook infer,” <https://research.facebook.com/downloads/infer/>, 2023.
- [48] “CodeQL Bounty Programs,” <https://securitylab.github.com/bounties/>.

- [49] J. Ruohonen, "An empirical analysis of vulnerabilities in python packages for web applications," in *2018 9th International Workshop on Empirical Software Engineering in Practice (IWESEP)*. IEEE, 2018, pp. 25–30.
- [50] M. Linares-Vásquez, G. Bavota, and C. Escobar-Velásquez, "An empirical study on android-related vulnerabilities," in *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*. IEEE, 2017, pp. 2–13.
- [51] D. Fried, A. Aghajanyan, J. Lin, S. Wang, E. Wallace, F. Shi, R. Zhong, W. Yih, L. Zettlemoyer, and M. Lewis, "InCoder: A generative model for code infilling and synthesis," *CoRR*, vol. abs/2204.05999, 2022.
- [52] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, "Evaluating large language models trained on code," *arXiv preprint arXiv:2107.03374*, 2021.
- [53] Z. Li, C. Wang, Z. Liu, H. Wang, S. Wang, and C. Gao, "Cctest: Testing and repairing code completion systems," *arXiv preprint arXiv:2208.08289*, 2022.
- [54] "ChatGPT," <https://openai.com/blog/chatgpt>.
- [55] "GPT4," <https://openai.com/research/gpt-4>.
- [56] Y. Li, D. Choi, J. Chung, N. Kushman, J. Schrittwieser, R. Leblond, T. Eccles, J. Keeling, F. Gimeno, A. Dal Lago *et al.*, "Competition-level code generation with alphacode," *Science*, vol. 378, no. 6624, pp. 1092–1097, 2022.
- [57] Z. Yu, W. Zheng, J. Wang, Q. Tang, S. Nie, and S. Wu, "Codecmr: Cross-modal retrieval for function-level binary source code matching," *Advances in Neural Information Processing Systems*, vol. 33, pp. 3872–3883, 2020.
- [58] Z. Li, P. Ma, H. Wang, S. Wang, Q. Tang, S. Nie, and S. Wu, "Unleashing the power of compiler intermediate representation to enhance neural program embeddings," in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 2253–2265.
- [59] E. Nijkamp, B. Pang, H. Hayashi, L. Tu, H. Wang, Y. Zhou, S. Savarese, and C. Xiong, "Codegen: An open large language model for code with multi-turn program synthesis," *arXiv preprint arXiv:2203.13474*, 2022.
- [60] "CodeGeeX," <https://models.aminer.cn/codegeex/blog/>.
- [61] R. Li, L. B. Allal, Y. Zi, N. Muennighoff, D. Kocetkov, C. Mou, M. Marone, C. Akiki, J. Li, J. Chim *et al.*, "StarCoder: may the source be with you!" *arXiv preprint arXiv:2305.06161*, 2023.
- [62] Z. Luo, C. Xu, P. Zhao, Q. Sun, X. Geng, W. Hu, C. Tao, J. Ma, Q. Lin, and D. Jiang, "WizardCoder: Empowering code large language models with evol-instruct," *CoRR*, vol. abs/2306.08568, 2023.
- [63] Q. Dong, L. Li, D. Dai, C. Zheng, Z. Wu, B. Chang, X. Sun, J. Xu, and Z. Sui, "A survey for in-context learning," *arXiv preprint arXiv:2301.00234*, 2022.
- [64] 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," vol. 33, 2020, pp. 1877–1901.
- [65] J. A. Prenner, H. Babii, and R. Robbes, "Can openai's codex fix bugs?: An evaluation on quixbugs," in *3rd IEEE/ACM International Workshop on Automated Program Repair, APR@ICSE 2022, Pittsburgh, PA, USA, May 19, 2022*. IEEE, 2022, pp. 69–75.
- [66] N. Nashid, M. Sintaha, and A. Mesbah, "Retrieval-based prompt selection for code-related few-shot learning," 2023.
- [67] Z. Li, C. Wang, P. Ma, C. Liu, S. Wang, D. Wu, and C. Gao, "On the feasibility of specialized ability stealing for large language code models," *CoRR*, vol. abs/2303.03012, 2023.
- [68] G. Poesia, A. Polozov, V. Le, A. Tiwari, G. Soares, C. Meek, and S. Gulwani, "SynchroMesh: Reliable code generation from pre-trained language models," in *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [69] S. Gao, X. Wen, C. Gao, W. Wang, and M. R. Lyu, "Constructing effective in-context demonstration for code intelligence tasks: An empirical study," *CoRR*, vol. abs/2304.07575, 2023.
- [70] M. Geng, S. Wang, D. Dong, H. Wang, G. Li, Z. Jin, X. Mao, and X. Liao, "Large language models are few-shot summarizers: Multi-intent comment generation via in-context learning," 2024.
- [71] T. Ahmed, K. S. Pai, P. T. Devanbu, and E. T. Barr, "Improving few-shot prompts with relevant static analysis products," *CoRR*, vol. abs/2304.06815, 2023.
- [72] H. Pearce, B. Tan, B. Ahmad, R. Karri, and B. Dolan-Gavitt, "Examining zero-shot vulnerability repair with large language models," *arXiv preprint arXiv:2112.02125*, 2021.
- [73] Google, "Google Bug hunter format," <https://bughunters.google.com/learn>, 2023.
- [74] WhiteSource, "Mend bolt," <https://www.mend.io/free-developer-tools/bolt>, 2022.
- [75] C. Wang, Y. Yang, C. Gao, Y. Peng, H. Zhang, and M. R. Lyu, "No more fine-tuning? an experimental evaluation of prompt tuning in code intelligence," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2022, pp. 382–394.
- [76] C. Bird, A. Bachmann, E. Aune, J. Duffy, A. Bernstein, V. Filkov, and P. Devanbu, "Fair and balanced? bias in bug-fix datasets," in *Proceedings of the 7th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*, 2009, pp. 121–130.
- [77] H. Wang, G. Ye, Z. Tang, S. H. Tan, S. Huang, D. Fang, Y. Feng, L. Bian, and Z. Wang, "Combining graph-based learning with automated data collection for code vulnerability detection," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 1943–1958, 2021.
- [78] P. Ma, Z. Li, A. Sun, and S. Wang, "'oops, did i just say that?'" testing and repairing unethical suggestions of large language models with suggest-critique-reflect process," *arXiv preprint arXiv:2305.02626*, 2023.
- [79] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy, "The pile: An 800gb dataset of diverse text for language modeling," *arXiv preprint arXiv: Arxiv-2101.00027*, 2020.
- [80] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [81] "Stanford Alpaca: An Instruction-following LLaMA model," <https://crfm.stanford.edu/2023/03/13/alpaca.html>.
- [82] S. Diao, R. Pan, H. Dong, K. Shum, J. Zhang, W. Xiong, and T. Zhang, "Lmflow: An extensible toolkit for finetuning and inference of large foundation models," <https://optimalscale.github.io/LMFlow/>, 2023.
- [83] A. Zeng, X. Liu, Z. Du, Z. Wang, H. Lai, M. Ding, Z. Yang, Y. Xu, W. Zheng, X. Xia, W. L. Tam, Z. Ma, Y. Xue, J. Zhai, W. Chen, Z. Liu, P. Zhang, Y. Dong, and J. Tang, "GLM-130b: An open bilingual pre-trained model," in *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [84] "vicuna," <https://vicuna.lmsys.org>.
- [85] "qian wen tongyi LLM," <https://qianwen.aliyun.com/>.
- [86] Z. Li, C. Wang, P. Ma, C. Liu, S. Wang, D. Wu, and C. Gao, "On the feasibility of specialized ability stealing for large language code models," *arXiv preprint arXiv:2303.03012*, 2023.
- [87] J. Wei, X. Wang, D. Schuurmans, M. Bosma, E. Chi, Q. Le, and D. Zhou, "Chain of thought prompting elicits reasoning in large language models," *arXiv preprint arXiv:2201.11903*, 2022.
- [88] J. L. Fleiss, "Measuring nominal scale agreement among many raters," *Psychological bulletin*, vol. 76, no. 5, p. 378, 1971.